

**IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE**

ARM LTD., a U.K. corporation,

Plaintiff,

v.

QUALCOMM INC., a Delaware corporation,
QUALCOMM TECHNOLOGIES, INC., a
Delaware corporation, and NUVIA, INC., a
Delaware corporation,

Defendants.

C.A. No. 22-1146-MN

PUBLIC REDACTED VERSION
(Filed July 22, 2024)

**DECLARATION OF NICHOLAS FUNG IN SUPPORT OF ARM LTD.'S
MOTIONS TO EXCLUDE AND STRIKE CERTAIN EXPERT OPINIONS OF MURALI
ANNAVARAM, PATRICK KENNEDY, JOHN COATES, AND JOEL STECKELARM**

VOLUME 2 OF 4 (EXHIBIT 6 – PART 1 OF 2)

OF COUNSEL:

YOUNG CONAWAY STARGATT &
TAYLOR, LLP

Daralyn J. Durie
Joyce Liou
MORRISON & FOERSTER LLP
425 Market Street
San Francisco, CA 94105
(415) 268-7000
ddurie@mofo.com
jliou@mofo.com

Erik J. Olson
MORRISON & FOERSTER LLP
755 Page Mill Road
Palo Alto, CA 94304
(650) 813-5600
ejolson@mofo.com

Kyle W.K. Mooney
Kyle D. Friedland
MORRISON & FOERSTER LLP
250 West 55th Street
New York, NY 10019
(212) 336-4092

Anne Shea Gaza (No. 4093)
Robert M. Vrana (No. 5666)
Samantha G. Wilson (No. 5816)
Rodney Square
1000 North King Street
Wilmington, DE 19801
(302) 571-6600
agaza@ycst.com
rvrana@ycst.com
swilson@ycst.com

Attorneys for Plaintiff Arm Ltd.

kmooney@mofo.com
kfriedland@mofo.com

Scott F. Llewellyn
MORRISON & FOERSTER LLP
4200 Republic Plaza
370 Seventeenth Street
Denver, CO 80202
(303) 592-2204
sllewellyn@mofo.com

Nicholas Rylan Fung
MORRISON & FOERSTER LLP
707 Wilshire Boulevard
Los Angeles, CA 90017
(213) 892-5348
nfung@mofo.com

Daniel P. Muino
MORRISON & FOERSTER LLP
2100 L Street, NW
Suite 900, Washington, D.C. 20037
(202) 887-1501
dmuino@mofo.com

**IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE**

ARM LTD., a U.K. corporation,

Plaintiff,

v.

QUALCOMM INC., a Delaware corporation,
QUALCOMM TECHNOLOGIES, INC., a
Delaware corporation, and NUVIA, INC., a
Delaware corporation,

Defendants.

C.A. No. 22-1146-MN

**HIGHLY CONFIDENTIAL –
ATTORENYS’ EYES ONLY – FILED
UNDER SEAL**

**DECLARATION OF NICHOLAS FUNG IN SUPPORT OF ARM LTD.’S
MOTIONS TO EXCLUDE AND STRIKE CERTAIN EXPERT OPINIONS OF MURALI
ANNAVARAM, PATRICK KENNEDY, JOHN COATES, AND JOEL STECKEL**

I, Nicholas Fung, declare as follows:

1. I am an attorney with the law firm of Morrison & Foerster LLP (“Morrison & Foerster”), counsel for Plaintiff Arm Ltd. (“Arm”) in the above-referenced action.
2. I submit this declaration in support of Arm’s Motions to Exclude the Expert Opinions of Murali Annavaram, Patrick Kennedy, John Coates, and Joel Steckel.
3. Attached hereto as **Exhibit 1** is a true and correct copy of the Opening Expert Report of Dr. Robert Colwell, dated December 20, 2023.
4. Attached hereto as **Exhibit 2** is a true and correct copy of the Opening Expert Report of Dr. Mike Chen, dated December 20, 2023.
5. Attached hereto as **Exhibit 3** is a true and correct copy of the Opening Expert Report of Mr. Todd Schoettelkotte, dated December 20, 2023.
6. Attached hereto as **Exhibit 4** is a true and correct copy of the Opening Expert Report of Mr. Guhan Subramanian, dated December 20, 2023.

7. Attached hereto as **Exhibit 5** is a true and correct copy of the Opening Expert Report of Dr. Ravi Dhar, dated December 20, 2023.

8. Attached hereto as **Exhibit 6** is a true and correct copy of the Opening Expert Report of Dr. Murali Annavaram, dated December 20, 2023.

9. Attached hereto as **Exhibit 7** is a true and correct copy of the Rebuttal Expert Report of Dr. Murali Annavaram, dated February 27, 2024.

10. Attached hereto as **Exhibit 8** is a true and correct copy of the Rebuttal Expert Report of Dr. Patrick Kennedy, dated February 27, 2024.

11. Attached hereto as **Exhibit 9** is an excerpt of a true and correct copy of the deposition transcript of John Coates, taken on April 19, 2024.

12. Attached hereto as **Exhibit 10** is a true and correct copy of the Expert Rebuttal Report of Mr. John Coates, dated February 27, 2024.

13. Attached hereto as **Exhibit 11** is a true and correct copy of the Expert Rebuttal Report of Dr. Joel Steckel, dated February 27, 2024.

14. Attached hereto as **Exhibit 12** is a true and correct copy of the Reply Expert Report of Dr. Murali Annavaram, dated March 25, 2024.

15. Attached hereto as **Exhibit 13** is a true and correct copy of the Opening Expert Report of Dr. Patrick Kennedy, dated May 20, 2024.

16. Attached hereto as **Exhibit 14** is an excerpt of a true and correct copy of the deposition transcript of Dr. Murali Annavaram, taken on June 27, 2024.

17. Attached hereto as **Exhibit 15** is an excerpt of a true and correct copy of the deposition transcript of Dr. Robert Colwell, taken on June 28, 2024.

18. Attached hereto as **Exhibit 16** is a true and correct copy of the Reply Expert Report of Mr. Todd Schoettelkotte, dated June 10, 2024.

19. Attached hereto as **Exhibit 17** is a true and correct copy of the Reply Expert Report of Mr. Guhan Subramanian, dated March 25, 2024.

20. Attached hereto as **Exhibit 18** is a true and correct copy of an email exchange from June 2021 and November 2022, with the subject line “[REDACTED],” produced by Defendants with Bates number QCARM_7434227.

21. Attached hereto as **Exhibit 19** is a true and correct copy of an email exchange from May 2021, with the subject line “[REDACTED],” produced by Defendants with Bates number QCARM_3535535.

22. Attached hereto as **Exhibit 20** is a true and correct copy of the Expert Reply Report of Dr. Patrick Kennedy, dated June 24, 2024.

23. Attached hereto as **Exhibit 21** is a true and correct copy of an email exchange from July 2021, with the subject line “Follow up,” produced by Arm with Bates numbers ARM_01305785 to - ARM_01305789.

I declare under penalty of perjury that the foregoing is true and correct to the best of my knowledge.

Executed this 10th day of July, 2024 at Los Angeles, California.

/s/ Nicholas Fung
Nicholas Fung

CERTIFICATE OF SERVICE

The undersigned hereby certifies that on July 10, 2024, a copy of the foregoing document was served on the counsel listed below in the manner indicated:

BY EMAIL/FTP

Jack B. Blumenfeld
Jennifer Ying
MORRIS, NICHOLS, ARSHT
& TUNNELL LLP
1201 North Market Street
P.O. Box 1347
Wilmington, DE 19899
jblumenfeld@morrisnichols.com
jying@morrisnichols.com

Isaac B. Zaur
Nora Niedzielski-Eichner
CLARICK GUERON REISBAUM LLP
220 Fifth Avenue, 14th Floor
New York, NY 10001
izaur@cgr-law.com
nniedzie@cgr-law.com

Catherine Nyarady
Anna R. Gressel
Madalyn G. Vaughn
Jacob A. Braly
Alexander M. Butwin
Samantha Mehring
PAUL, WEISS, RIFKIND,
WHARTON & GARRISON LLP
1285 Avenue of the Americas
New York, NY 10019
cnyarady@paulweiss.com
agressel@paulweiss.com
mvaughn@paulweiss.com
jbraly@paulweiss.com
abutwin@paulweiss.com
smehring@paulweiss.com

Karen L. Dunn
William A. Isaacson
Melissa F. Zappala
Erin J. Morgan
Brian Shiue
Anna P. Lipin
PAUL, WEISS, RIFKIND,
WHARTON & GARRISON LLP
2001 K Street, NW
Washington, DC 20006
kdunn@paulweiss.com
wisaacson@paulweiss.com
mzappala@paulweiss.com
ejmorgan@paulweiss.com
bshiue@paulweiss.com
alipin@paulweiss.com

Andrea L. D'Ambra
Susana Medeiros
Kira Latham
NORTON ROSE FULBRIGHT US LLP
1301 Avenue of the Americas
New York, NY 10019
andrea.dambra@nortonrosefulbright.com
susana.medeiros@nortonrosefulbright.com
kira.latham@nortonrosefulbright.com

YOUNG CONAWAY STARGATT &
TAYLOR, LLP

/s/ Anne Shea Gaza

Anne Shea Gaza (No. 4093)

Robert M. Vrana (No. 5666)

Samantha G. Wilson (No. 5816)

Rodney Square

1000 North King Street

Wilmington, DE 19801

(302) 571-6600

agaza@ycst.com

rvrana@ycst.com

swilson@ycst.com

Attorneys for Plaintiff Arm Ltd.

Exhibit 6

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

ARM LTD.,

Plaintiff,

v.

QUALCOMM INC., QUALCOMM
TECHNOLOGIES, INC. and NUVIA, INC.,

Defendants

C.A. No. 22-1146 (MN)

OPENING EXPERT REPORT OF DR. MURALI ANNAVARAM

December 20, 2023

HIGHLY CONFIDENTIAL – SOURCE CODE - ATTORNEY’S EYES ONLY

TABLE OF CONTENTS

[illegible]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

II. INTRODUCTION

1. My name is Murali Annavaram. I have been retained as an expert in this action on behalf of Defendants Qualcomm Inc., Qualcomm Technologies (collectively, “Qualcomm”), and Nuvia, Inc. (“Nuvia”) (together, “Defendants”). I am being compensated for my work on this case at my standard consulting rate of \$600 per hour. I am also being reimbursed for expenses that I may incur. My compensation is not contingent upon the results of my analysis or the substance of my opinions or testimony.

A. Summary of Opinions

2. For this report, I have been asked to analyze the materials discussed in this report and offer my expert opinions on the reasonableness of Qualcomm’s process to identify and remove Nuvia-sourced ARM RTL and related test and debug code¹ from certain Qualcomm codebases. For reasons described further below, I will refer to this process as the “Swap Out.” [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

3. I submit this report to describe my opinions related to the Swap Out. For background, I provide a technology overview, including Register-Transfer Level code (RTL) and how processors and system-on-chips (“SoCs”) are developed and represented in RTL. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

¹ The source code discussed in this Report is predominantly RTL (e.g. file extensions, .v, .sv or .svh). Even though there are instances where the source code is not what would be classified as RTL, I will generally reference the source code in this Report as RTL unless otherwise noted.

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

4. After ARM terminated Nuvia's license, it is my understanding that Qualcomm undertook efforts to identify and remove all Nuvia-sourced ARM RTL from Qualcomm Codebases. My opinion, based on my expertise in processor design, is that Qualcomm designed and implemented a thorough process to identify Nuvia-Sourced ARM RTL in the Qualcomm Codebases based on my review of the materials discussed below, [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

B. Reservations

5. I expect to be called to provide expert testimony regarding opinions formed resulting from my analysis of the issues considered in this report if asked about those issues by the

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

court or by the private parties' attorneys. If called to testify, I may use demonstratives to explain the concepts and issues discussed below. I may also use various documents produced in this case that refer to or relate to the matters discussed in this report to present my opinions. I may also discuss my own work, teaching, and publications in the field, and knowledge of the state of the art in the relevant time period. I may rely on handbooks, textbooks, technical literature, my own personal experience in the field, and other relevant materials and/or information to demonstrate the state of the art in the relevant period and the evolution of relevant technologies.

6. I reserve the right to modify or supplement my opinions, as well as the basis for my opinions, after considering new positions set forth by ARM LTD. ("ARM"). For example, I may update my opinions based on additional opinions that ARM's experts may present and information I may receive in the future or additional work I may perform in connection with these opinions and information.

7. It is my understanding that discovery is still ongoing. For example, I understand that the parties are still taking depositions, and completing discovery, which may result in updating written discovery. I reserve the right to modify or supplement my opinions, as well as the basis for my opinions, based on any documents, testimony, or other evidence that may emerge during the course of this matter.

8. It is also my understanding that ARM may submit an expert report corresponding to this report. I reserve the right to rebut any positions taken in that report.

III. QUALIFICATIONS AND BACKGROUND

9. My curriculum vitae ("CV") is attached as Appendix A and provides a summary of my background, education, and professional experience.

10. I am an expert in the field of SoC, CPU and GPU architecture and microarchitecture, mobile systems, and datacenter computing. My research focuses on designing

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

energy efficient and high performance chip multiprocessors (CMPs), server computer systems, mobile SoCs, microarchitectural innovations and instruction set architecture (ISA) enhancements, RTL based modeling and simulations of SoC, CPU and GPU microarchitectural blocks, verification and testing of hardware designs, software based simulation of SoC and processor components to perform wide range design space tradeoff analysis, compiler schemes for translating high level language applications into efficient instructions mapped to the underlying processor ISA, interconnection network architectures, reliability of computing platforms, and superconducting computer architecture. I have done many of these studies both in industry and academia.

11. I have been a faculty member at the Ming-Hsieh Department of Electrical and Computer Engineering, with a joint appointment in the Computer Science department at the University of Southern California since 2007. I currently hold the Lloyd Hunt Chair Professorship at USC. In the past I held the position of Dean's Professor until June 2023, and Robert G. and Mary G. Lane Early Career Chair until Aug 2017. I also held the Rukmini Gopalakrishnachar visiting chair at the Indian Institute of Science.

12. I co-authored Parallel Computer Organization and Design, a widely-used textbook for graduate-level computer architecture courses, which addresses the design of modern high-performance CPUs and prepares students for a career designing the computer systems of the future. I teach several undergraduate and graduate computer architecture courses at USC, including topics such as CPU and GPUs microarchitecture, server and cloud computing, mobile and edge computing systems, systems for machine learning, and memory system design. These classes have included EE109 – Introduction to Embedded Systems, EE557 – Advanced Computer Architecture, EE653 – Advanced Topics in Microarchitecture, and other special topic courses.

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

13. Prior to joining USC, I spent nearly 7 years at Intel as a senior research staff member. At Intel, I worked on a broad range of topics including processor microarchitecture techniques to improve power-performance tradeoffs, efficient coherence and consistency models for chip multiprocessors (CMPs), 3D stacked memory technologies, energy per instruction (EPI) throttling schemes, database performance optimizations on Intel architectures, and the impact of process variations on CPU's timing analysis. I spent a year at Nokia as a visiting scientist working on mobile SoCs, including studying the tradeoffs between compute, sensing and communication in SoCs.

14. I received my Bachelor of Technology degree in Computer Science from the National Institute of Technology in Warangal, India, in 1993, and my Master of Science degree in Computer Science and Engineering from Colorado State University in 1996. I received my Ph.D. degree in Computer Science and Engineering from The University of Michigan in 2001.

15. I have published over 100 well-cited research papers with a total citation count exceeding 8600. My work on CPU design using 3D die stacking, energy efficient computing on CMPs through EPI throttling, and energy efficient sensing and computing in SoCs have each been cited over 300 times.

16. I have received numerous awards including, for example, the IBM Faculty Partnership award in 2008, the Best Paper Award at the 2009 IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS), and the National Science Foundation CAREER award in 2010.

17. I was named an IEEE fellow for my contributions to heterogeneous architectures for energy-efficient computing systems. I am a Distinguished Member of the Association of Computing Machinery (ACM).

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

18. My research has been published at the major computing conferences in the field of computer architecture, including the International Symposium on Computer Architecture (ISCA), the High-Performance Computer Architecture (HPCA), and the Symposium on Microarchitecture (MICRO).

19. My research on server energy efficiency has been selected for the IEEE Micro TopPicks award for the most influential computer architecture paper of the year 2013. I have given keynote presentations at the Swedish Multicore Symposium in 2018 and at the IEEE International On-Line Test Symposium in 2014. My group's research won Best Paper and Best Student paper finalist nominations at the Supercomputing 2019 conference. I was the Frontiers in Technology Distinguished Speaker at UC Merced during 2022.

20. For my research and publications in this field, I was inducted into the ACM Special Interest Group on Microarchitecture Hall of Fame in December 2015. I am one of only 47 researchers over the past 48 years at that time to receive this honor. I was also inducted into the IEEE Computer Society Technical Committee on Computer Architecture's High Performance Computer Architecture Hall of Fame. I am one of only 42 researchers over the past 22 years at that time to receive this honor. Finally, I was inducted into the ACM International Symposium on Computer Architecture (ISCA) Hall of Fame in June 2017. I am one of only 81 researchers to have received this honor over the past 44 years at the time of the award.

21. I served as the General Co-Chair for ISCA 2018 and served as the Technical Program Chair for HPCA 2021. I am going to serve as a Technical Program Chair of the International Conference on Supercomputing, 2024. I have also served as a program committee member for several international conferences and symposiums, including ISCA, MICRO, HPCA, SIGMETRICS, DSN, HiPC, and ASPLOS.

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

22. I served as an associate editor for the technical journal of the Association for Computer Machinery Transactions on Design Automation of Electronic Systems (ACM TODAES), and I served as a journal reviewer for the ACM Transactions on Computer Systems (ACM TOCS), ACM Transactions on Embedded Computing Systems (ACM TECS), ACM Computing Surveys, ACM Transactions on Architecture and Code Optimization (ACM TACO), and IEEE Top Picks in Microarchitecture 2010.

23. I am an inventor or co-inventor of several issued patents.

24. I have graduated 15 PhD students at USC, and 11 of them were solely advised by me. Of these, 6 are now in academia (both in the USA and abroad) as professors working on various computer architecture related problems.

25. My Research Lab at USC is nicknamed SCIP (Super Computing In Pocket). SCIP research areas include: energy efficiency through heterogeneous computing; reliability of high-performance computing; bandwidth efficient big data computing; runtime systems design to enable dispersed computing; hardware-assisted secure and private machine learning; and building innovative sensor data collection platforms to improve human performance. SCIP engages with researchers in industry and academia to tackle compelling computer system challenges.

26. I am also the inaugural director of the USC-Meta Center for Research and Education in AI and Learning (REAL@USC-Meta Center). REAL@USC advances foundations for cooperative algorithmic optimization, hardware innovations for AI, and AI education accessibility.

27. My research group has published more than 30 papers on the topic of CPU and GPU microarchitecture alone at top-tier computer architecture conferences in the past 10 years, where the typical acceptance rates are less than 20%. In particular, my group published 3 papers

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

at the flagship computer architecture conference ISCA during 2016, which is a record for the highest number of published papers by a single research group.

28. My research group has designed and evaluated the energy efficient techniques for mobile systems, including the KNOWME wireless body area sensor networks, energy efficient management of sensor subsystems on mobile phones, and privacy preserving traffic monitoring from mobile devices.

29. My group has done extensive experimentation and analysis of our proposed research ideas using a wide range of instruction set architecture enhancements, microarchitecture block designs of CPUs and GPUs, operating system modifications, system-level hardware modifications, and design enhancements. We routinely work with RTL modules, instruction set architectures (ISAs), test and verification tools, software and RTL simulation infrastructures.

IV. BASIS FOR OPINIONS

30. My qualifications are summarized in Section III of this report. My full curriculum vitae is attached as Appendix A to this report.

31. As part of my preparation for writing this report, I reviewed the materials listed in Appendix B to this report. These materials include, but are not limited to, the following: spreadsheets and scripts responsible for identifying Nuvia-sourced ARM RTL or other related code, scripts for the removal of related information, deposition transcripts, and copies of the documents described in this report.

32. I also reviewed RTL in Qualcomm Codebases that was produced in this litigation. For that code review, I accessed codebases for multiple snapshots of RTL during my time reviewing code at the Prosearch source code review location in Los Angeles, CA. I provide further description of this code in my analysis below.

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

33. I further had several discussions with Qualcomm engineers Nick Jones (Director of Engineering, Qualcomm) and Bob Pflederer (Senior Director, Qualcomm) regarding the Swap Out, who were both substantially involved in the process.

V. TECHNOLOGY BACKGROUND

34. In what follows, I will provide an overview of processor technology and RTL representation used in microarchitectures.

A. Processor Technology

35. A System-on-Chip or SoC design places a variety of components such as processors, caches, memories, and input/output devices all on a single piece of silicon. Such a design offers both reduced latency for executing an application and reduced power consumption than a design in which the various components are placed on separate chips. One reason for improved SoC power efficiency and performance is that the components on the SoC can communicate with each other using on-chip wires that have lower resistance than off-chip pins. SoCs also enable integration of heterogeneous technologies such as central processing unit (CPU), graphics processing unit (GPU), volatile memory such as DRAM that holds application programs when they are executing, and non-volatile storage such as Flash memory that holds applications and other media files even when the device is turned off.

36. One of the components typically included on a SoC is one or more Central Processing Unit (CPU) cores. Each CPU core may include components or blocks such as a data cache, an instruction cache, an Arithmetic Logic Unit (ALU), a Memory Management Unit (MMU), and a Floating Point Unit (FPU). A CPU may include only a single CPU core (called a single-core CPU or single-core processor) or may include multiple CPU cores (called a chip multiprocessor (CMP) or multi-core processor).

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

37. A CPU is designed to perform operations on data. Designers implement these operations in circuitry using a collection of techniques such as microarchitecture design, the placement of gates and routing of wires across the chip, selection of process technology nodes, and selection of the cell implementation libraries. This circuitry is represented in a type of computer language referred to as RTL code, which I will describe further in the next section.

B. RTL

38. RTL describes the implementation of the CPU's microarchitecture. RTL design is a digital design methodology that focuses on the transfer of data between registers within a digital system. It serves as an abstraction level between the high-level behavioral description of a system and its physical implementation in hardware. At the RTL level, designers describe the functionality and behavior of the system in terms of registers, data flow, and control signals.

39. A designer may develop an RTL description of a digital circuit manually using a Hardware Description Language (HDL) such as Verilog or automatically from a higher-level language using an RTL synthesis tool. After the RTL representation is completed, it is then transformed through a series of hardware compilers, place and route tools and cell libraries to create representation that helps with the fabrication of the actual physical device.

40. The RTL description may be organized into groups of code referred to as "modules," with each module likewise organized into smaller groups of code referred to as "submodules." Organizing the RTL into modules and submodules improves the efficiency in managing large projects by allowing code to be modularized and re-used in different aspects of the projects. For example, the RTL modules for complex functions implemented in a CPU can be treated as abstract boxes with input and output pins. Hence different module developers can easily interface without knowing the details of the module design.

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

41. Verilog has some commonality with other high level programming languages like C or Fortran. For example, a module is, conceptually, very similar to a subroutine. However, the behavior of a Verilog program is generally dictated by the way signals are routed across connected modules rather than by the order of code that appears within the Verilog program.

42. For a simple example of RTL, consider a D (Delay) flip flop that transfers an input value to an output on the falling edge of a clock signal. This is a common component in digital logic circuits, like CPUs. The flip flop can be used to temporarily store data during CPU operation. The behavior of such a component could be described with RTL as follows:

```
Module Dff(q, clock, data);
  output    q;
  reg       q;
  input      clock, data;

  initial
    q=0;

  always
    @(negedge clock) #1 q=data;
endmodule
```

43. The “always” statement in that RTL description indicates that when the clock goes from high to low (i.e., on a falling or negative edge) the output register “q” is assigned the input value “data” after a delay of one-time unit.

44. The use of RTL in the design of a CPU provides numerous advantages including ease of development of the processor using a high-level programming abstraction, much like writing software using C/C++ to accomplish an application task. The hardware designers do not need to worry too much about how the digital gates, such as AND and NOR gates, are implemented in a particular process technology node.

45. A second advantage provided by RTL is design validation in which the designers are able to simulate and test their design using commercial tools. Such testing allows the designers

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

to identify and correct problems much earlier in the design process and with much less cost than would be incurred later during the fabrication process.

46. A third advantage provided by RTL is modularity. Designers can create libraries of RTL components that can be easily reused across an entire design. For example, the flip flop described in the example RTL code above may be a defined module that is re-used throughout CPU and SoC design. These reusable blocks can be easily integrated into new designs, promoting design reuse, reducing development time, and enabling faster prototyping and system assembly.

VI. BACKGROUND FACTS

47. In this section, I summarize my review of information relevant to my analysis, including a background of the early development work Nuvia did for the [REDACTED]

[REDACTED]

[REDACTED] and ARM RTL made available through the ARM Connect platform.

A. Nuvia

48. Nuvia was founded in 2019 by Gerard Williams (CEO, who formerly worked at Apple), Manu Gulati (SVP Silicon Engineering, who formerly worked at Google), and John Bruno (SVP Engineering, who formerly worked at Google) with the goal of bringing higher energy efficiency and performance to the server market. QCARM_0002749 at 59-60; QCARM_2414840. Nuvia focused on the design and development of chips used in data centers, by specializing the chip by prioritizing efficiency for operations common in data centers and less common in desktops, laptops, and mobile devices. [REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

[REDACTED]

49. Qualcomm acquired Nuvia on March 15, 2021.³

B. Qualcomm

50. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

51. I have been informed that during the early development of [REDACTED] at Nuvia, Nuvia entered into license agreements with ARM that enabled Nuvia to access certain ARM IP, including some RTL modules (ARM RTL). Nuvia instantiated some ARM RTL in the [REDACTED] SoC by downloading the ARM RTL through ARM Connect (using Nuvia's credentials), which is a website made available to ARM licensees such as Nuvia. I provide additional information regarding ARM Connect below.

52. I have been informed that ARM terminated the license agreement with Nuvia on or around March 2022, about one year after Qualcomm completed its acquisition of Nuvia. [REDACTED]

[REDACTED]

³ <https://www.qualcomm.com/news/releases/2021/03/qualcomm-completes-acquisition-nuvia>

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

[REDACTED]

53. I have been informed that Qualcomm had separate license agreements with ARM, which were in effect at the time that Qualcomm completed the acquisition of Nuvia and continues to remain in effect at the time that I provide this Report. Through those licenses, Qualcomm has a license for and corresponding access to at least the same ARM RTL as downloaded under the Nuvia license. A licensee like Qualcomm can obtain the ARM RTL through ARM Connect by using its credentials supplied by ARM as part of a license with ARM. I will refer throughout my report to the ARM RTL that is available on ARM Connect as “Nuvia-sourced ARM RTL” for copies of ARM RTL that Nuvia downloaded from ARM Connect using Nuvia’s credentials, and “Qualcomm-sourced ARM RTL” for copies of ARM RTL that Qualcomm downloaded from ARM Connect using Qualcomm’s credentials.

54. As I will describe in more detail below, Qualcomm undertook reasonable efforts to identify Nuvia-sourced ARM RTL and completely remove it from its source code repositories. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

C. ARM Connect

55. [REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

[REDACTED]

[REDACTED]

56. ARM allows its customers to download ARM IP such as RTL, documentation, or release notes through its Product Download Hub, formerly known as Connect.⁴ I will focus on Connect as it was the platform available during the relevant time period discussed in this report, namely 2019 through 2022.⁵

57. As described in the Connect user guide: [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

⁴ <https://developer.arm.com/documentation/107572/latest/>.

⁵ According to ARM documentation, Connect was closed around August 2022, and transitioned to the Product Download Hub that appears to contain similar functionality.
<https://developer.arm.com/documentation/107571/latest/>.

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

Connect Product Page for Cortex-M3

Id. at 74.

58. Within the Connect platform, the download process includes [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

VII. ANALYSIS

59. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

A. The Qualcomm Codebases

60. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

61.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

62.

[REDACTED]

[REDACTED]

[REDACTED]

63.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

[REDACTED]

64. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

65. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

[REDACTED]

66. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

67. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

68. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

69. [REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

[REDACTED]

70. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

71. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

72.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

B.

[REDACTED]

73.

[REDACTED]

[REDACTED]

1.

[REDACTED]

74.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

75. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[illegible]

2. Scripts for Identification

81. To aid in the methodical identification of ARM RTL,

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

82.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED]

[REDACTED]

83.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] [REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

[REDACTED]

85.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

86.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

3. [REDACTED]

87. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

88.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

89.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

i. [REDACTED]

92. [REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

(1) [REDACTED]

94. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

95.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

i.

[REDACTED]

96.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

97. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

ii. Regular Expressions

98.

[REDACTED]

[REDACTED]

[REDACTED]

99.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

- [REDACTED]

- [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

- [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

- [REDACTED]

[REDACTED]

[REDACTED]

iii. [REDACTED]

103. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

104.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

iv.

[REDACTED]

105.

[REDACTED]


[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY



106.


[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY



v. Metadata Modification Methods

107.

[Redacted text]

[Redacted text]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

[REDACTED]

[REDACTED]

108. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

109. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

(2) [REDACTED] [REDACTED]

110. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] [REDACTED]

[REDACTED]

[REDACTED]

i. [REDACTED]

[REDACTED] [REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]


114.

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY



115.

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

ii. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[illegible]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]

[REDACTED]

[REDACTED]

122. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

123.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

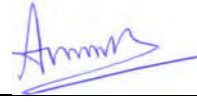
[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

I certify under penalty of perjury that the foregoing is true and correct.

Date: December 20, 2023



Murali Annavaram, Ph.D.

Los Angeles, California

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

ARM LTD.,

Plaintiff,

v.

QUALCOMM INC., QUALCOMM
TECHNOLOGIES, INC. and NUVIA, INC.,

Defendants

C.A. No. 22-1146 (MN)

OPENING EXPERT REPORT OF DR. MURALI ANNAVARAM

APPENDIX

A

Murali Annavaram

Lloyd Hunt Chair Professor
Ming Hsieh Department of Electrical & Computer Engineering
Thomas Lord Department of Computer Science
University of Southern California
Los Angeles, CA 90089-2562
TEL: (213) 740-3299
annavara@usc.edu
<http://scip-lab.usc.edu>



Professional Experience

- 7/2023-Current: **Lloyd Hunt Chair in Electrical and Computer Engineering**, Ming Hsieh Electrical and Computer Engineering Department, Thomas Lord Department of Computer Science (courtesy), University of Southern California
- 8/2022-7/2023: **Rukmini Gopalakrishnachar Chair Visiting Professor**, Indian Institute of Sciences, Bangalore, India.
- 4/2022- 7/2023t: **Dean's Professor**, Ming Hsieh Electrical and Computer Engineering Department, Thomas Lord Department of Computer Science, University of Southern California
- 12/2017- Current: **Professor**, Ming Hsieh Electrical and Computer Engineering Department, Thomas Lord Department of Computer Science, University of Southern California
- 4/2012- 12/2017: **Assoc. Professor**, Ming Hsieh Electrical and Computer Engineering Department, Thomas Lord Department of Computer Science, University of Southern California
- 8/2007- 4/2012: **Asst. Professor**, Ming Hsieh Electrical and Computer Engineering Department, Thomas Lord Department of Computer Science (courtesy), University of Southern California

- 4/2010-12/2017: **Robert G. and Mary G. Lane Early Career Chair Professor**, Ming Hsieh Electrical and Computer Engineering Department, University of Southern California
- 2/2020-2/2021: Consulting Faculty Scientist, Facebook
- 8/2014-12/2014: Visiting faculty on Sabbatical, Caltech
- 2/2007-8/2007: Faculty Researcher, Nokia Research Labs
- 8/2001-2/2007: Senior Research Scientist, Intel
- 5/1995-8/1995: Summer Intern, AT&T Bell Labs.
- 9/1993-7/1994: Software Engineer, Wipro Systems, India.

Education

- The University of Michigan, Ann Arbor, MI; 1996-2001
Ph.D. Computer Science and Engineering
(Dissertation Title: Prefetch Mechanisms that Acquire and Exploit Application Specific Knowledge; Advisor: Prof. Edward Davidson)
- Colorado State University, Fort Collins, CO; 1994-1996
M.S, Computer Science and Engineering
- National Institute of Technology, Warangal, India; 1989-1993
B.Tech, Computer Science

Research Summary

Murali Annavaram is the Lloyd Hunt Chair Professor in the Ming-Hsieh Department of Electrical and Computer Engineering at the University of Southern California. He also holds a joint appointment with the Thomas Lord department of Computer Science. He held the Rukmini Gopalakrishnchar Visiting Chair Professorship at the Indian Institute of Science, Bangalore. He is the founding Director of the USC-Meta Center for Research and Education in AI and Learning. He is also the architecture thrust leader for the DISCoVER NSF Expeditions in Computing Center at USC. For his numerous publications, he was inducted into the hall of fame at three top-tier computer architecture conference venues, ISCA, HPCA and MICRO. He held the Dean's Chair and Robert G. and Mary G. Lane Early Career Chair professorships at USC in the past.

His research focuses on:

- ◇ Superconductive computing
- ◇ Energy efficiency through heterogeneous computing
- ◇ Reliability of high performance computing
- ◇ Bandwidth efficient big data computing,
- ◇ Runtime systems design to enable dispersed computing
- ◇ Hardware-assisted secure and private machine learning
- ◇ Building innovative sensor data collection platforms to improve human performance.

Prior to his appointment at USC he worked at Intel Microprocessor Research Labs from 2001 to 2007 working on energy efficient system design and die-stacking

architectures. In 2007 he was a visiting researcher at the Nokia Research Center, Palo Alto working on mobile phone-based wireless traffic sensing using virtual trip lines.

Murali received the NSF CAREER award in 2010 and an IBM Faculty Partnership award in 2008, Facebook faculty award in 2019. He received the Stevens Institute's Innovation Inside curriculum award for jointly developing a mobile systems design course.

Murali co-authored Parallel Computer Organization and Design, a widely used textbook to teach both the basic and advanced principles of computer architecture. This book is used in graduate computer architecture courses in several reputed universities around the world.

Murali received the Ph.D. degree in Computer Engineering from the University of Michigan, Ann Arbor, in 2001. He is a Senior Member of ACM. He is an IEEE Fellow

Honors and Awards

- 12/2023: ACM Distinguished Member
- 7/2023: Lloyd Hunt Chair in Electrical Power
- 8/2022: Rukmini Gopalakrishnachar Chair Visiting Professorship, Indian Institute of Science
- 4/2022: Dean's Professorship in ECE and CS departments
- 4/2022: Architecture thrust leader of the Discover Expeditions grant from NSF (largest computing grant given at the NSF CISE division)
- 3/2022: Frontiers in Technology Distinguished Speaker, UC Merced
- 12/ 2021: Founding and Inaugural Director, USC-Meta Center for Research and Education in AI and Learning
- 11/2021: IEEE Fellow
- 11/2021: Keynote speaker at the King Fahd University of Petroleum & Minerals annual KIXX forum
- 10/2021: NeurIPS 2021 Outstanding Reviewer Award
- 12/2020: Best paper award for FedML.AI infrastructure
- 11/2019: Best paper and Best Student paper finalist nominations, Supercomputing 2019 conference
- 11/2018: Keynote speaker for Swedish Multicore Symposium
- 6/2017: Inducted into the ACM ISCA hall of fame for publishing eight papers at the ISCA conference (one of 81 researchers over the last 44 years)
- 3/2016: Patent on improving mobile system's battery efficiency licensed
- 12/2015: Inducted into the ACM SIGMICRO hall of fame for publishing eight papers at the MICRO conferences (one of 47 researchers over the last 48 years)
- 7/2014: Keynote speaker at the 2014 IEEE International On-Line Test Symposium, Spain.
- 6/2013: IEEE Micro Top Picks Award for the paper titled "KnightShift-Scaling Energy Proportionality Wall Through Server-level Heterogeneity."
- 8/2011: Holder of Robert G. and Mary G. Lane Early Career Chair 2011-Present
- 7/2010: NSF CAREER award 2010

- 4/2012: Best Paper Nomination at the 2012 IEEE International Symposium on Workload Characterization (IISWC) for his work “Wireless Body Area Networks: Where Does the Energy Go?”
- 10/2009: Best Paper Award at the 2009 IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS) for his joint work on “Optimal Time-Resource Allocation for Activity-Detection via Multimodal Sensing.”
- 10/2009: Winner of the 2009 Body Computing Award for project titled The KNOWME Network: Energy-Efficient Activity-Detection for Pediatric Obesity
- 8/2008: Senior Member of Association of Computing Machinery (ACM)
- 8/2008: Senior Member of Institute of Electrical and Electronics Engineers (IEEE)
- 8/2008: IBM Faculty Partnership Award 2008.
- 12/2007: USC Stevens Institute’s Innovation Inside Curriculum Award for developing EE579 course on mobile application development at USC
- 3/2007: Nokia Exemplary Team Achievement award (highest award at the Nokia research labs)
- 5/2002, 5/2005: Two Intel divisional recognition awards (highest awards at the corporate technology division at Intel).
- 8/1996-8/1997: University of Michigan EECS Department Fellowship.
- 1989: 32nd Rank in AP Engineering and Medical Common Entrance Test.
- 1987: 14th Rank in State-wide High School Graduation Exam, India.
- 1987-1993: National Merit Scholarship, India.

Consulting Experience

- June 2014-Oct 2015. Testifying expert in the damages phase of **Warf-vs-Apple**. Testified in a jury trial. Deposed by the defendant’s counsel.
- Nov 2015-May 2017. ITC matter on various haptic patents **Immersion-vs-Apple**. Delivered expert report and witness statements.
- Aug 2015-May 2017. **Futurelink-vs-Intel** ITC case related to various patents regarding server designs. Delivered expert report and a supplemental expert report and was deposed.
- Oct 2016-June 2018. **Qualcomm-vs-Apple** ITC case on a patent associated with GPUs. Delivered expert report and a supplemental expert report on infringement and validity of the patent. Deposed by the opposing counsel. Testified at ITC.
- June 2018-Feb 2019. **Qualcomm-vs-Apple** district case on a patent associated with GPUs. Delivered expert report and a supplemental expert report on infringement and validity of the patent. Deposed by the opposing counsel. Testified at jury trial.
- Feb 2018-Dec 2018. **NVIDIA-vs-ZiiLabs** ITC case on a patent associated with GPUs. Delivered expert report and a supplemental expert report on non-infringement and invalidity of the patent. Deposed by the opposing counsel.
- Sept 2019-Dec 2022. **PACT-vs-Intel** IPR proceedings. Submitted expert declaration.

- Dec 2019-June 2023. **XMTT**-vs-Intel IPR proceedings. Submitted expert declaration and deposed.
- Dec 2019-June 2023. **XMTT**-vs-Intel District case on a patent. Submitted expert report and deposed.
- June 2020-Oct 2020. **Samsung**-vs-KIPB. District case on a patent. Submitted expert report and deposed.
- Aug 2020-Dec 2020. **Samsung**-vs-Arbor. District case on a patent. Submitted expert report and deposed.
- Sept 2019-Current. **VLSI**-vs-Intel case on a patent dispute. Expert report submitted and deposed. Testified at trials in Texas.
- June 2021-June 2022. **Netlist**-vs-Google case on a patent dispute.
- June 2021-April 2023. **Williams**-vs-Apple case on a trade secret dispute.
- July 2022-Current. **Rivos**-vs-Apple case on a trade secret dispute.
- May 2023-Current, **Qualcomm**-vs-ARM case on a licensing dispute.

Research Projects axnd Support

Annaram's research is supported through the following peer-reviewed research grants and industrial research gifts.

Active:

1. VMWARE, "Enhancing Privacy Through Hardware Trusted Execution Environments," M. Annaram (PI, \$265K, July 2022-July 2025)
2. NSF, "SHF:Small:ML Accelerator Cohort Architecture," M. Annaram (PI, 600K), June 2022-June 2025.
3. NSF, Expeditions in Computing, "DISCOVER: Design and Integration of Superconductive Computation for Ventures beyond Exascale Realization," M. Annaram (co-PI, ~300K/year), May 2022-Aug 2029. Total grant amount ~15M total funds.
4. USC-Meta Center for Research and Education in AI and Learning. M. Annaram (Director), Dec 2021-Dec 2026 (~5M total funds).
5. DARPA FastNICs: "DIAMOND: Distributed Training of Massive Models at Bandwidth Frontiers," M. Annaram (Co-PI ~750K), May 1st 2020-May 1st 2024, Total Grant Amount \$2,062,540.

Completed:

1. Avast/Borsat/Intel Gift, "Secure and Privacy Preserving Machine Learning: Foundations and Scalable System Design", M. Annaram (Co-PI ~50K/year), Nov 2020 – Nov 2023.
2. IARPA SuperTools: "ColdFlux: CAD Methodologies and Tools for Single Flux Quantum Based Superconductive Electronics," M. Annaram (Co-PI, ~\$500,000), M. Pedram (PI). May 1st 2017- Dec 31st 2022, Total Grant Amount ~\$13M.

3. NSF, "CNS – 2002874/200282: Collaborative Research: MLWiNS: A Coding-Centric Approach to Robust, Secure, and Private Distributed Learning over Wireless," M. Annavaram (Co-PI~100K), Aug 1st 2020-Dec 31st 2021, Total Grant Amount (~300K)
4. DARPA Dispersed Computing Grant: "ApaC: Adaptive Pricing and Coding for Dispersed Computing," M. Annavaram (Co-PI -33%, ~\$1,406,782), S. Avestimehr (PI-33%), B. Krishnamachari (Co-PI-33%), April 1st 2017- April 1st 2021, Total Grant Amount \$ 4,217,754.
5. Facebook Gift, "Machine Inference at the Edge," M. Annavaram (PI), May 1st 2020-April 30th 2021, Total Gift Amount 100,000.
6. NSF, "SHF:Small Semantically Aware Graph Storage and Processing," M. Annavaram (PI-75%), A. Ortega (Co-PI, 25%), Aug 15th 2017 – Aug 14th 2021, Total Grant Amount \$400,000.
7. Facebook Gift, "Privacy Preserving Machine Inference," M. Annavaram, Nov 1st 2019-Nov 1st 2020, Total Gift Amount 50,000.
8. Samsung: "Energy Efficient Graph Processing on SSDs," M. Annavaram (PI), , June 15th 2017 – June 14th 2018, Total Grant Amount \$100,000.
9. IARPA: "TILES: Tracking Individual Performance with Sensors," M. Annavaram (Co-PI, ~\$500,000), S. Narayanan (PI). May 1st 2017- Oct 31st 2020, Total Grant Amount ~\$12M.
10. Ming-Hsieh Institute Mini Grant, "Creating the USC Technology Enabled Aging Mind Center (TEAMc)", M. Annavaram (co-PI), U. Mitra (PI), \$8,996.
11. NSF CAREER Grant# 0954211, "From Nonstop-Monitoring to Nano-ISA: An Adaptive Multi-Dimensional Framework for Processor Reliability", M. Annavaram (PI), Funded, August 1, 2010-July 31, 2017, \$427,664.
12. DARPA PERFECT Grant: "Embedded POWER Optimized Systems Using Near and Super-threshold Computing Fabric (EMPOWER)," M. Annavaram (Co-PI -33%), M. Pedram (PI-33%), K. Roy (Co-PI-33%), Phase 3 funded Aug 1st 2015-Oct 1st 2016, Total Grant Amount \$2,155,102.
13. NSF Grant# 1219186, "SHF:Small: Benchmarking of Transient and Intermittent Errors and Their Application to Microarchitecture", M. Annavaram (PI-50%), M. Dubois (Co-PI-50%), Sept 1, 2012 – August 31 2017, \$400,000.
14. DARPA PERFECT Grant: "Low-power and Error-resilient Digital Components Realized in Deeply-scaled CMOS (LEDRA)," M. Annavaram (Co-PI -0%), M. Fritze (PI), Funded Phase 1 for ~\$1.5M Dec 1st 2012-Dec 1st 2014, Total Grant Amount: ~\$8M. Grant money returned due to publication restrictions.
15. DARPA IRIS Grant: "Techniques for Estimating Reliability in COTS Ics (TERCI)," M. Annavaram (Co-PI -10%), March 1st 2013-March 1st 2014, \$125,000.
16. NSF Grant# 0954211, "IEEE International Symposium on Workload Characterization (IISWC) Student Subsidy Proposal", M. Annavaram (PI), Jan 1, 2011-Dec 31, 2011, \$5,000.
17. NSF Grant# 0831545, "A Game Theoretic Framework for Privacy Preservation in Community-Based Mobile Applications", M. Annavaram (PI-50%), B. Krishnamachari (Co-PI-50%), October 1, 2008 – September 30, 2011, \$249,999.

18. NSF REU Supplement Grant# 0939625, “A Game Theoretic Framework for Privacy Preservation in Community-Based Mobile Applications”, M. Annavaram (PI-50%), B. Krishnamachari (Co-PI-50%), June 1, 2009 – September 30, 2011, \$16,000.
19. NSF Grant# 0834798, “A Holistic Design Approach to Reliability Using 3D Stacked Monitoring”, M. Annavaram (PI), Sept 1, 2008 – August 31 2011, \$402,904.
20. NSF REU Supplement Grant# 0939623, “A Holistic Design Approach to Reliability Using 3D Stacked Monitoring”, M. Annavaram (PI), July 7, 2009 – August 31 2011, \$16,000.
21. NSF Grant# 0834799, “Trade-offs Between Static Power, Performance and Reliability in Future Chip Multiprocessors”, M. Annavaram (PI-50%), M. Dubois (Co-PI-50%), September 1, 2008 – August 31, 2010, \$200,000.
22. NIH/National Center for Minority Health and Health Disparities, Grant# 53-4503-6037, “Mobile device biomonitoring to prevent and treat obesity in underserved minority youth (KnowMe Study)” M. Annavaram (Principal-30%), D. Metz (Principal-30%), U.Mitra (Co-PI-10%), S. Narayanan (Co-PI-10%), G. Sukhatme (Co-PI-10%), N. Medvidovic (Co-Principal-10%), May 1, 2008 – April 30, 2010, \$948,348.
23. Qualcomm Research Grant, “Mobile Metabolic Health Monitoring”, M. Annavaram (Co-PI -10%), D. Metz (Co-PI-10%), U.Mitra (Co-PI-10%), S. Narayanan (PI-50%), G. Sukhatme (Co-PI-10%), August 1, 2008 – July 31, 2009, \$50,000.
24. Oracle Equipment Grant. Virtex-5 OpenSPARC Research Platform Boards, August 2011, \$16,000.
25. Qualcomm Equipment Grant. Snapdragon Mobile Development Platform Boards, MSM8660, August 2011, \$13,500.
26. Nokia Equipment Grant. 50 Lumia900 Phones, Jan 2013, \$40,000.
27. Annenberg Micro Seminar series grant. Aug 2011, \$2,500.
28. Rose Hills Foundation Science and Engineering Fellowship grant. June 2011, \$2,000.
29. Ming-Hsieh Institute Mini Grant, “The Smartphone In Your Pocket – Refreshing The Hardware Of EE579”, M. Annavaram, Jan 2011, \$8,996.
30. Nokia Research Gift, “Mobile Energy Efficiency”, M. Annavaram (PI), September 1, 2007, \$50,000.
31. Nokia Research Gift, “Mobile Systems for Health”, M. Annavaram (PI), May 1, 2008, \$50,000.
32. Nokia Research Gift, “Mobile Systems for Health”, M. Annavaram (PI), April 1, 2009, \$50,000.
33. IBM Faculty Award Researc Gift, “3D Stacking for Reliability”, M. Annavaram (PI), August 2008, \$25,000.
34. USC Startup Package, August 15th 2007, \$70,000 Equipment, \$50,000 Research Expenses, \$120,000 RA Support, 4 Summer Months Salary

Student Advising

Current Phd Students
1. Haipeng Zha
2. Tingting Tan
3. Keshav Balasubramanian
4. Abdulla Alshabanah
5. Yongqin Wang
6. Rachit Rajat
7. Tara Renduchintala
8. James Flemings
9. Lei Gao
10. Chaoyi Jiang
11. Jonghyun Lee
12. Hossein Entezari Zarch

Phd Students Graduated	Graduation Date	First Employment
1. Kimish Patel (co-advised with Prof. Pedram)	Sept 2010	Nvidia
2. Yi Wang (co-advised with Prof. Krishnamachari)	March 2011	Juniper
3. Jinho Suh (co-advised with Prof. Dubois)	Jan 2012	Intel
4. Bardia Zandian	Nov 2012	Qualcomm
5. Waleed Dweik	Dec 2014	Asst. Professor, University of Jordan, Amman
6. Daniel Wong	June 2015	Asst. Professor, UC, Riverside
7. Hyeran Jeon	June 2015	Asst. Professor, San Jose State University Currently Asst Professor @ UC Merced
8. Sangwon Lee	Dec 2015	VMWare
9. Mohammad Abdel-Majeed	Feb 2016	Asst. Professor, University of Jordan, Amman
10. Qiumin Xu	May 2017	Google
11. Abdulaziz Tabbakh	Jan 2018	Asst. Professor, KAUST, Saudi Arabia.

12. Gunjae Koo	May 2018	Asst. Professor, Hongkik Univ, Korea Currently Asst Professor @ Korea University
13. Kiran Matam	Dec 2019	Facebook
14. Krishna Narra	Aug 2020	Google
15. Hanieh Hashemi	Aug 2023	Apple

MS Students Advised	Graduation Date	First Employment
1. Sabyasachi Ghosh	June 2012	Riverbed Networks
2. Suk Hun Kang	June 2010	Samsung
3. Kumar Dabbiru	June 2012	Google
4. Melina Demertzi	June 2013	Oracle
5. Zhifeng Lin	Aug 2019	Google

BS Students Mentored	Graduation Date	First Employment
1. Thomas Punihaole	June 2010	UCLA, MS
2. Willy Long	June 2013	Duke, PhD
3. Justin Kuang	June 2014	Stanford, MS
4. Zhifeng Lin	June 2013	USC, PhD
5. Francisco Romero	June 2015	Stanford, PhD
6. Julia Chen	June 2014	Oracle

Phd Student Thesis Committee Member
1. Richard Hankins
2. Weirong Jiang
3. Hua Liu
4. Jianwei Chen
5. Mehrtash Manochchri

Teaching and Course Development at USC

- EE579: Wireless and Mobile Networks Design and Laboratory. Developed this course from ground-up to focus on challenges associated with mobile

phone design challenges and usage models. The course covers issues related to energy efficiency in mobile technologies and data privacy issues.

- EE599: Energy Efficiency and Reliability in Information and Communication Technologies. Developed this new experimental course focusing on topics related to energy proportional computing, datacenter power usage efficiency, server reliability and methods to tradeoff energy efficiency and reliability.
- EE357, EE457 and EE557: Enhanced computer architecture course sequence to bring parallel architecture and software development issues into mainstream architecture courses.

Textbook Co-Authorship

- Parallel Computer Organization and Design. Michel Dubois, Murali Annavaram and Per Stenström. Cambridge Press. Published in August 2012.

Conference Organizational Responsibilities:

- 12/2021: Technical program chair of IEEE 34th International Symposium on Computer Architecture and High-Performance Computing
- 2/2020: Technical program chair of HPCA 2021
- 6/2018: General co-chair of ISCA 2018
- 4/2016: Track Chair: Green Computing, Design Automation and Test in Europe (DATE), 2016.
- 6/2014: Vice chair computer architecture track, High Performance Computing (HiPC), 2014

Other Professional Activities and Service

- Associate Editor
 - Journal of Parallel and Distributed Computing
 - ACM Transactions on Design Automation of Electronic Systems (2011-2014)
- Journal Reviewer
 - ACM Transactions on Computers
 - ACM Transactions on Embedded Computing Systems
 - ACM Computing Surveys
 - ACM Transactions on Architecture and Code Optimization
 - IEEE Top Picks in Microarchitecture 2010, 2015, 2016, 2017.
- Program committee member (including external review committee memberships)
 - IEEE/ACM Architectural Support for Programming Languages and Operating Systems (ASPLOS) 2018, 2019, 2020.
 - NeurIPS 2020, 2022
 - Sigmetrics 2016, 2017, 2018, 2019, 2020, 2021, 2022, 2023
 - International Symposium on Computer Architecture (ISCA), 2015, 2016

- Dependable Systems and Networks (DSN), 2015, 2020
- Design Automation and Test in Europe (DATE), 2015, 2016
- IEEE/ACM International Symposium on Microarchitecture (MICRO) 2007, 2011, 2014, 2015, 2016, 2018, 2022, 203
- IEEE High Performance Computer Architecture (HPCA) 2009, 2010, 2014, 2017, 2022, 2023
- International Parallel and Distributed Processing Systems (IPDPS) 2014, 2015, 2017, 2018.
- IEEE International Symposium on Performance Analysis of Software and Systems (ISPASS) 2011, 2012, 2018
- International Conference on Mobile Computing, Applications, and Services (MobiCASE) 2011
- IEEE Workshop of Energy Efficient Design (WEED) 2011, 2012
- IEEE International Symposium on Workload Characterization (IISWC) 2010, 2012, 2014, 2015, 2016
- IEEE/ACM International Conference on Parallel Architecture and Compilation Techniques (PACT) 2010
- International Conference on Mobile Computing, Applications, and Services (MobiCASE) 2009, 2010, 2011
- Program Track chair
 - Workshop on Exascale Evaluation and Research Techniques (EXERT) 2010, 2011
 - Workshop on Duplicating, Deconstructing, and Debunking (WDDD) 2011, 2012, 2013. 2014, 2015
- Local arrangements chair HPCA 2004, Registration chair for HPCA 2005 and HPCA 2006, Registration chair for PACT 2010, Finance Chair for IISWC 2010.
- Reviewer for IEEE Micro Special Issue on Hardware/Software Interactions 2008, MICRO 2006, MICRO 2005, ISCA 2005, HPCA 2004, ISCA 2003, HPCA 2001, MICRO 2001, ASPLOS 2000, ISCA 2000, PACT 1999, ISCA 1998, MICRO 1998, MICRO 1997, HPCA 1997.
- Multiple NSF proposal review panels

Graduate Student Experience

- Research Assistant: *Jan 1998 - Aug 2001, The University of Michigan.*
 - I worked as a research assistant under Prof. Edward Davidson working on my Ph.D. thesis. My dissertation research focused on reducing cache miss stalls of irregular applications such as databases by hardware and software assisted prefetching.
- Teaching Assistant: *Sept 1996 – Dec 1997, The University of Michigan.*
 - Teaching assistant for a programming and introductory data structures course (one semester) and a graduate level operating systems course (two semesters). Job responsibilities included leading weekly class discussions, designing homework for the operating systems class and grading homework and exams.

- Research Assistant: *May 1996 - Aug 1996, Colorado State University*
 - I worked as a research assistant under Prof. Walid Najjar in the Pebbles research group at Colorado State University. I enhanced a Machine Independent Dataflow Code (MIDC) format that is used by a Sisal compiler for automatic generation of blocking threads from Sisal programs. I also enhanced a cycle accurate multithreaded machine simulator, called ID, to execute the blocking threads. This simulation infrastructure is used for evaluating the memory system performance of multithreaded machines (more details are available in conference paper [16]).
- Teaching Assistant: *August 1994 – May 1996, Colorado State University*
 - Teaching assistant for introductory personal computing course (two semesters), introductory programming language course (one semester) and automata theory course (one semester). Job responsibilities included leading weekly class discussions, lab in-charge and grading homework and exams.

Refereed Publications

Short note on Conference Publications: Here I am paraphrasing information from three relevant articles (1) “Evaluating Computer Scientists and Engineers For Promotion and Tenure,” by Prof. David Patterson, *et Al.* (2) “Committee on Academic Careers for Experimental Computer Scientists,” by National Research Council (3) “Judging the Impact of Conference and Journal Publications in Computer Architecture,” University of California, San Diego Computer Engineering Research Guidelines. In my research area, primarily Computer Architecture, our top tier publication venues are conferences where the conference paper quality meets or exceeds journal publication. Conferences in our field exceed journals in selectivity, visibility, circulation, number of reviews, expertise of reviewers and program committee members. Our conferences have about 5 reviewers and each submitted paper is 20-25 pages in length. Our conference papers represent high quality finished research as they report full evaluations with detailed simulations or models. The program committee consists of about 25 world-renowned researchers and the acceptance rates are less than 30%; I listed acceptance rates where I was able to get the data. Hence, I will first list the refereed full conference papers with acceptance rates wherever available, followed by journals related to my inter-disciplinary work and finally workshop papers. In the list below * represents authors are my PhD students either advised solely by me or co-advised. Underline represents the primary senior author. If multiple primary senior authors they are all underlined in the author list.

Conference Papers

1. [MICRO] H. Zha*, S. Tannu and M. Annavaram. SuperBP: Design Space Exploration for SFQ Branch Predictors. In *proceedings of 55th IEEE/ACM International Symposium on Microarchitecture*, October 2022. (Acceptance rate 101/424, 24%).

2. [ISCA] Y. Wang*, R. Rajat*, and **M. Annavaram**. LAORAM: Look Ahead ORAM for Recommendation Model Privacy. . In *proceedings of the International Symposium on Computer Architecture (ISCA)*, June 2023 (Acceptance rate 79/373, 21%).
3. [MICRO] R. Rajat*, Y. Wang*, and **M. Annavaram**. PageORAM: An Efficient DRAM Page Aware ORAM. In *proceedings of 55th IEEE/ACM International Symposium on Microarchitecture*, October 2022. (Acceptance rate 83/369, 22%).
4. [NAACL] E.S. Markowitz, K. Balasubramanian*, M. Mirtaheeri, **M. Annavaram**, **A. Galstyan**, **G.V. Steeg**. StATIK: Structure and Text for Inductive Knowledge Graph Completion. In *2022 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. July 2022.
5. [ISPASS] Y. Wang*, E. Suh, W. Xiong, **M. Annavaram** and **H. Lee**. Characterization of MPC-based Private Inferences for Transformer-based Models. In *Proceedings of the International Symposium on Performance Analysis of Systems and Software*, May 2022.
6. [ICASSP] T. Feng, H. Hashemi*, **M. Annavaram** and **S. Narayanan**. Enhancing Privacy Through Domain Adaptive Noise Injection for Speech Emotion Recognition. In *proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*, May 2022.
7. [IPDPS] T. Tang*, H. Hashemi*, R. Ali, S. Avestimehr, and **M. Annavaram**. AVCC: Adaptative Verifiable Coded Computing. In *proceedings of International Conference on Parallel and Distributed Processing Systems*, May 2022 (Acceptance rate 46/474 10% Round1 acceptance).
8. [AAAI] E. Ceyani, C. He, K. Balasubramanian, **M. Annavaram**, and **S. Avestimehr**. SpreadGNN: : Serverless Multi-task Federated Learning for Molecular Graphs. In *proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI-22)*, Feb 2022. (Acceptance rate 1349/9251, 15%).
9. [HPCA] H. Zha*, N. Katam, M. Pedram and **M. Annavaram**. HiPerRF: A Dual-Bit Dense Storage SFQ Register File. In *proceedings of the International Conference on High-Performance Computer Architecture (HPCA)*, Feb 2022. (Acceptance rate 54/260, 21%).
10. [NSDI] A. Eisenman, K. Matam, S. Ingram, D. Mudigere, R. Krishnamoorthi, K. Nair, M. Smelyanskiy and **M. Annavaram**. Check-N-Run: A Checkpointing System for Training Deep Learning Recommendation Models. In *Proceedings of the 19th USENIX Symposium on Networked Systems Design and Implementation (NSDI)* April 2022.
11. [MICRO] H. Hashemi, Y. Wang, and **M. Annavaram**. DarKnight: An Accelerated Framework for Privacy and Integrity Preserving Deep Learning Using Trusted Hardware. In *proceedings of 54th IEEE/ACM International Symposium on Microarchitecture*, October 2021. (Acceptance rate 94/423, 22%).
12. [RECSYS] K. Balasubramanian*, A. Alsabnah*, J. Choe* and **M. Annavaram**. cDLRM: Look Ahead Caching for Scalable Training of Recommendation Models. In *proceedings of 15th ACM Conference on Recommender Systems*, Oct 2021 (Acceptance rate 49/267, 18%).
13. [CLOUD] K. Narra*, Z.Lin*, Y. Wang*, K. Balasubramanian*, and **M. Annavaram**. Origami Inference: Private Inference Using Hardware Enclaves. In

- proceedings of IEEE International Conference on Cloud Computing (Short paper)*, Sept 2021 (Acceptance rate 23.8%).
14. [IPDPS] K. Matam, H. Hashemi, and **M. Annavaram**. MultiLogVC: Efficient Out-of-Core Graph Processing Framework on Flash Storage. In *proceedings of International Conference on Parallel and Distributed Processing Systems*, May 2021.
 15. [NeurIPS] C. He, **M. Annavaram**, and S. Avestimehr. Group Knowledge Transfer: Federated Learning of Large CNNs at the Edge. In *Advances in Neural Information Processing Systems*. (Acceptance rate 1900/9454, 20%). Dec 2020.
 16. [ICDCS] K. Narra, Z. Lin, S. Avestimehr, G. Ananthanarayanan, and **M. Annavaram**. Collage Inference: Using Coded Redundancy for Lowering Latency Variation in Distributed Image Classification Systems/ In *proceedings of International Conference on Distributed Computing Systems*, July 2020 (Acceptance rate 105/584, 18%).
 17. [SC] K. Narra*, Z. Lin, M. Kiamari, S. Avestimehr, and **M. Annavaram**. Distributed Matrix Multiplication Using Speed Adaptive Coding. In *proceedings of SuperComputing (SC)*, Nov 2019 (Acceptance rate 72/344, 21%). **Best Paper and Best Student Paper Finalist Nomination.**
 18. [ICS] Q. Xu*, H. Naghibijouybari, S. Wang, N. Abu-Ghazaleh, and **M. Annavaram**. GPUGuard: Mitigating Contention Based Side and Covert Channel Attacks on GPUs In *proceedings of the International Symposium on SuperComputing (ICS)*, June 2019 (Acceptance rate 45/167, 27%).
 19. [ISCA] Y. Oh, G. Koo*, **M. Annavaram**, W. Ro. Linebacker: Preserving Victim Cache Lines in Idle Register Files of GPUs. In *proceedings of the International Symposium on Computer Architecture (ISCA)*, June 2019 (Acceptance rate 62/365, 17%).
 20. [ISCA] K. Matam*, G. Koo*, H. Zha*, H. Tseng, and **M. Annavaram**. GraphSSD: Graph Semantics Aware SSD. In *proceedings of the International Symposium on Computer Architecture (ISCA)*, June 2019 (Acceptance rate 62/365, 17%).
 21. [NIPS] M. Yu, Z. Lin*, K. Narra*, S. Li, Y. Li, N. Kim, A. Schwing, **M. Annavaram**, and S. Avestimehr. GradiVeQ: Vector Quantization for Bandwidth-Efficient Gradient Aggregation in Distributed CNN Training. In *Advances in Neural Information Processing Systems*, pp. 5129-5139. 2018. (Acceptance rate 1011/4856, 21%).
 22. [IPDPS] G. Koo*, H. Jeon, Z. Liu, N. Kim and **M. Annavaram**. CTA-Aware Prefetching and Scheduling. In *proceedings of the Symposium on Parallel Distributed Computing*, May, 2018 (Acceptance rate 38/461, 8%).
 23. [HPCA] A. Tabbakh*, X. Qian and **M. Annavaram**. G-TSC: Timestamp Based Coherence for GPUs. In *proceedings of the International Conference on High-Performance Computer Architecture (HPCA)*, Feb 2018. (Acceptance rate 54/260, 21%).
 24. [MICRO] G. Koo*, K. Matam*, Te I, K. Narra*, Jing Li, **H. Tseng**, S. Swanson, and **M. Annavaram**. Flash-Summarizer: Trading Communication with Computing Near Storage. In *proceedings of the IEEE International Symposium on Microarchitecture*, Oct 2017. (Acceptance rate 61/325, 19%)

25. [ISCA] G. Koo*, Y. Oh, W. Ro, and M. Annavaram. Access Pattern-Aware Cache Management for Improving Data Utilization in GPU. In *proceedings of the International Symposium on Computer Architecture (ISCA)*, June 2017 (Acceptance rate 54/322, 17%).
26. [IPDPS] A. Tabbakh*, X. Qian and M. Annavaram. Power Efficient Sharing-Aware GPU Data Management. In *proceedings of the Symposium on Parallel Distributed Computing*, June 2017.
27. [IPDPS] B. Li, J. Sun, M. Annavaram and N.S. Kim. Elastic-Cache: GPU Cache Architecture for Efficient Fine- and Coarse-Grained Cache-Line Management. In *proceedings of the Symposium on Parallel Distributed Computing*, June 2017.
28. [HPCA] M. Abdel-Majeed*, H. Jeon*, A.S. Bejestan, M. Pedram, and Murali Annavaram. Pilot Register File: Energy Efficient Register File for GPUs. In *proceedings of the International Conference on High-Performance Computer Architecture (HPCA)*, Feb 2017. (Acceptance rate 50/224, 22%)
29. [HPCA] Z. Liu, S. Gilani, Murali Annavaram, and N.S. Kim. G-Scalar: Cost-effective generalized scalar execution architecture for power-efficient GPUs. In *proceedings of the International Conference on High-Performance Computer Architecture (HPCA)*, Feb 2017. (Acceptance rate 50/224, 22%)
30. [ISCA] Q. Xu*, H. Jeon*, K. Kim*, W. Ro, M. Annavaram. Efficient Intra-SM Slicing through Dynamic Resource Partitioning for GPU Multiprogramming. In *proceedings of the International Symposium on Computer Architecture (ISCA)*, June 2016. (Acceptance rate 57/291, 20%)
31. [ISCA] M. Yoon*, S. Lee*, K. Kim*, G. Koo*, W. Ro and M. Annavaram. Virtual Thread: Maximizing Thread-Level Parallelism beyond GPU Scheduling Limit. In *proceedings of the International Symposium on Computer Architecture (ISCA)*, June 2016. (Acceptance rate 57/291, 20%)
32. [ISCA] Y. Oh, K. Kim*, M. Yoon*, J. Park, Y. Park, W. Ro, and M. Annavaram. APRES: Improving Cache Efficiency by Exploiting Load Characteristics on GPUs. In *proceedings of the International Symposium on Computer Architecture (ISCA)*, June 2016. (Acceptance rate 57/291, 20%)
33. [ICS] M. Abdel-Majeed*, D. Wong*, J. Kuang and M. Annavaram. Origami: Folding Warps for Energy Efficient GPUs. In *proceedings of the International conference on Supercomputing (ICS)*, June 2016. (Acceptance rate 43/183, 23%)
34. [HPCA] D. Wong*, N. Kim and M. Annavaram. Warped-Approximation: Using Value Similarity for approximate computing on GPUs. In *proceedings of the International Conference on High-Performance Computer Architecture (HPCA)*, March 2016. (Acceptance rate 53/240, 23%)
35. [HPCA] K. Kim*, S. Lee, M. Yoon, G. Koo, W. Ro and M. Annavaram. Warped-Preexecution: A GPU Preexecution Approach to Improve Latency Hiding. In *proceedings of the International Conference on High-Performance Computer Architecture (HPCA)*, March 2016. (Acceptance rate 53/240, 23%)
36. [MICRO] H. Jeon*, G. Ravi, N.S. Kim and M. Annavaram. GPU register file virtualization. In *proceedings of the IEEE International Symposium on Microarchitecture*, Dec 2015. (Acceptance rate 61/283, 22%)

37. [IISWC] G. Koo*, H. Jeon* and **M. Annavaram**. Revealing Critical Loads and Hidden Data Locality in GPGPU applications. In proceedings of the *IEEE International Symposium on Workload Characterization (IISWC)*, October 2015.
38. [IISWC] D. Wong*, J. Chen* and **M. Annavaram**. Traveling to the Edge of Energy Proportionality. In proceedings of the *IEEE International Symposium on Workload Characterization (IISWC)*, October 2015.
39. [DSN] M. Abdel-Majeed*, W. Dweik*, H. Jeon* and **M. Annavaram**. Warped-RE: Low-Cost Fault Detection and Correction in GPUs. In proceedings of the *International Conference on Dependable Systems and Networks (DSN)*, June 2015.
40. [ISCA] S. Kim, K. Kim, G. Koo, H. Jeon, W. Ro, and **M. Annavaram**. Warped-Compression: Enabling Power Efficient GPUs Through Register Register Compression. In proceedings of the *International Symposium on Computer Architecture (ISCA)*, June 2015.
41. [ITC] H. Jeon*, G. Loh and **M. Annavaram**. “RAS Support for Wide-I/O Stacked DRAM”. In proceedings of the *IEEE International Test Conference (ITC)*, Nov 2014.
42. [IISWC] Q. Xu*, H. Jeon*, and **M. Annavaram**. “Characterization and analysis of GPGPU design for Graph Applications”. In proceedings of the *IEEE International Symposium on Workload Characterization (IISWC)*, Oct 2014.
43. [PACT] Q. Xu* and **M. Annavaram**. “PATs: Pattern Aware Scheduling and Power Gating for GPGPUs”. In proceedings of the *International Conference on Parallel Architecture and Compilation Techniques (PACT)*, Aug 2014.
44. [DSN] W. Dweik*, M. Abdel-Majeed* and **M. Annavaram**. “Warp-Shield: Tolerating Hard Faults in GPGPUs”. In proceedings of the *International Conference on Dependable Systems and Networks (DSN)*, June 2014.
45. [GOMAC] W. Dweik* and **M. Annavaram**. SignTest: Signature-based Adaptive Periodic Testing. In *Government Microcircuit Applications and Critical Technology Conference, (GOMACTech)*, April 2014.
46. [GOMAC] W. Dweik*, M. Dubois and **M. Annavaram**. Reliability-Aware Exceptions: Tolerating Intermittent Faults in Microprocessor Array Structures. In *Government Microcircuit Applications and Critical Technology Conference, (GOMACTech)*, April 2014.
47. [DATE] W. Dweik*, M. Dubois and **M. Annavaram**. Reliability-Aware Exceptions: Tolerating Intermittent Faults in Microprocessor Array Structures. In proceedings of the *Design Automation and Test in Europe (DATE)*, March 2014.
48. [HPCA] D. Wong *, and **M. Annavaram**. Implications of High Energy Proportional Servers on Cluster-wide Energy Proportionality. In *proceedings of the International Conference on High-Performance Computer Architecture (HPCA)*, Feb 2014.
49. [MICRO] M. Abdel-Majeed*, D. Wong* and **M. Annavaram**. Warped Gates: Gating Aware Scheduling and Power Gating for GPGPUs. In *proceedings of the International Symposium on Microarchitecture (MICRO)*, Dec 2013.
50. [DSN] J. Suh*, M. Dubois and **M. Annavaram**. PHYS: Profiled-Hybrid Sampling for Soft Error Reliability Benchmarking. In proceedings of the

International Conference on Dependable Systems and Networks (DSN), June 2013.

51. [HPCA] M. Abdel-Majeed*, and **M. Annavaram**. Warped Register File: A Power Efficient Register File for GPGPUs. In *proceedings of the International Conference on High-Performance Computer Architecture (HPCA)*, Feb 2013.
52. [MICRO] H. Jeon* and **M. Annavaram**. Warped-DMR: Light-weight Error Detection for GPGPUs. In *proceedings of the International Symposium on Microarchitecture (MICRO)*, Dec 2012.
53. [MICRO] D. Wong* and **M. Annavaram**. Scaling the energy proportionality wall through server-level heterogeneity. In *proceedings of the International Symposium on Microarchitecture (MICRO)*, Dec 2012.
54. [IISWC] S. Lee* and **M. Annavaram**. Wireless Body Area Networks: Where Does the Energy Go?. In *proceedings of the IEEE International Symposium on Workload Characterization (IISWC)*, Nov 2012. **(Best Paper Award Nominee)**
55. [IISWC] M. Demertzi*, B.Zandian*, R. Rojas* and **M. Annavaram**. Benchmarking Instruction Set Architecture for Non-Transient Errors. In *proceedings of the IEEE International Symposium on Workload Characterization (IISWC)*, Nov 2012.
56. [DSN] B. Zandian* and **M. Annavaram**. Software-based Infield Wearout Monitoring for Synchronous Digital Chips. In *proceedings of the International Conference on Dependable Systems and Networks (DSN) – Fast Abstracts*, June 2012.
57. [DSN] W. Dweik* and **M. Annavaram**. Signature-based Online Periodic Fault Tolerance for Microprocessors. In *proceedings of the International Conference on Dependable Systems and Networks (DSN) – Fast Abstracts*, June 2012.
58. [SECON] Y.Wang*, **B. Krishnamachari**, and **M. Annavaram**. Semi-Markov State Estimation and Policy Optimization for Energy Efficient Mobile Sensing. In *proceedings of 9th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, 2012.
59. [ISQED] M.Abdel-Majeed*, S.Chen and **M. Annavaram**. A case for 3D stacked analog circuits in high-speed sensing systems. In *proceedings of the IEEE International Symposium on Quality Electronics Design (ISQED)*, March 2012.
60. [HPCA] J. Suh*, **M. Annavaram** and M. Dubois. MACAU: A Markovian Model for Reliability Evaluation of Caches Under SBUs and MBUs. In *proceedings of the International Symposium on High Performance Computer Architecture (HPCA)*, Feb 2012 (Acceptance rate 36/210, 17%).
61. [IISWC] M.Demertzi*, **M. Annavaram** and M. Hall. Analyzing the effect of compiler optimizations on application reliability. In *proceedings of the 2011 IEEE International Symposium on Workload Characterization (IISWC)*, Nov 2011.
62. [EMBC] S. Kim, M. Li, S. Lee*, U. Mitra, A. Emken, D. Spruijt-Metz, **M. Annavaram**, S. Narayanan. High-level Descriptions of Real-Life Physical Activities Using Latent Topic Modeling of Multimodal Sensor Signals. In *proceedings of the 33rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, Aug 2011.
63. [DSN] B.Zandian* and **M. Annavaram**. Cross Layer Resilience Using Wearout Aware Design Flow. In *proceedings of the International Conference on*

- Dependable Systems and Networks (DSN)*, pages 279-290, June 2011 (Acceptance rate 26/148, 18%).
64. [ISCA] M. Manocheri*, **M. Annavaram** M. Dubois. CPCC: Correctable Parity Protected Cache. In *proceedings of the International Symposium on Computer Architecture (ISCA)*, pages 223-234, June 2011 (Acceptance rate 40/208, 19%).
 65. [SIGMETRICS] J. Suh*, M. Manocheri, **M. Annavaram** and M. Dubois. Soft Error Benchmarking of L2 Caches with PARMA. In *proceedings of the International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, pages 85-96, June 2011 (Acceptance rate 26/177, 15%).
 66. [HPCA] N. Madan*, A. Buyuktosunoglu, P. Bose and **M. Annavaram**. Guarded Power Gating in a Multi-Core Setting. In *proceedings of the International Symposium on High Performance Computer Architecture (HPCA)*, pages 291-300, Feb 2011 (Acceptance rate 46/238, 19%).
 67. [MICRO] J. Chen, K. Dabbiru*, D. Wong*, **M. Annavaram** and M. Dubois. Adaptive and Speculative Slack Simulations of CMPs on CMPs. In *proceedings of the International Symposium on Microarchitecture (Micro)*, pages 523-534, Dec 2010. (Acceptance rate: 45/245, 18%)
 68. [DSN] B. Zandian*, W. Dweik*, S. Kang*, T. Punihaole* and **M. Annavaram**. Continuous Reliability Monitoring Using Adaptive Critical Path Testing. In *proceedings of the International Conference on Dependable Systems and Networks (DSN)*, pages 151-160, July 2010 (Acceptance rate: 39/170, 23%)
 69. [IPSN] Y.Wang*, B. Krishnamachari, Q. Zhao and **M. Annavaram**. Markov-Optimal Sensing Policy for User State Estimation in Mobile Devices. In *proceedings of 9th International Conference on Information Processing in Sensor Networks 2010 - IP Track (IPSN-2010)*, pages 268-278, Apr 2010 (Acceptance rate: 20/117, 17%).
 70. [ITC] T-Y. Hsieh, M. Breuer, **M. Annavaram**, S. Gupta, K-J. Lee. Tolerance of Performance-degrading Faults for Effective Yield Improvement. In *Proceedings of the 40th Annual International Test Conference (ITC)*, Pages 1-10, Nov 2009.
 71. [MOBICASE] Y.Wang*, B. Krishnamachari, Q. Zhao and **M. Annavaram**. Towards the Tradeoff between Energy Efficiency and User State Estimation Accuracy in Mobile Sensing. In *Proceedings of the International Conference on Mobile Computing, Applications, and Services (MOBICASE)*, Oct 2009.
 72. [EMBC] G. Thatte, M. Li, A. Emken, U. Mitra, S. Narayanan, **M. Annavaram** and D. Spruijt-Metz. Energy-Efficient Multihypothesis Activity-Detection for Health-Monitoring Applications. In *Proceedings of the 31st Annual International Conference of IEEE Engineering in Medicine and Biology (EMBC09)*, September 2009.
 73. [ICPP] J. Chen, **M. Annavaram** and M. Dubois. Exploiting Simulation Slack to Improve Parallel Simulation Speed. In *Proceedings of the 2009 International Conference on Parallel Processing (ICPP)*, pages 371-378, Sept 2009. (Acceptance Rate: 30%)
 74. [DCOSS] G. Thatte, V. Rozgic, M. Li, S. Ghosh*, U. Mitra, S. Narayanan, **M. Annavaram**, and D. Spruijt-Metz. Optimal allocation of time-resources for multi hypothesis activity-level detection. In *Proceedings of the 5th IEEE International*

Conference on Distributed Computing in Sensor Systems (DCOSS'09), June 2009.
(Best Paper Award)

75. **[MOBISYS]** Y. Wang*, J. Lin, **M. Annavaram**, Q. A. Jacobson, J. Hong, B. Krishnamachari and N. Sadeh. A Framework of Energy Efficient Mobile Sensing for Automatic User State Recognition. In *proceedings of 7th Annual International Conference on Mobile Systems, Applications and Services (MobiSys)*, pages 179-192, June 2009 (Acceptance Rate, 26/128, 20%).
76. **[SIGMETRICS]** J. Chen, **M. Annavaram**, M. Dubois. Slacksim: : A Platform for Parallel Simulations of CMPs on CMPs. In *proceedings of International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, June 2009. (Acceptance Rate: 26%)
77. **[MSWIM]** H. Liu, B. Krishnamachari and **M. Annavaram**. Game theoretic approach to location sharing with privacy in a community-based mobile safety application. In *Proceedings of the 11th international Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems MSWiM '08*. Pages 229-238, October 2008.
78. **[MOBISYS]** B. Hoh, M. Gruteser, R. Herring, J. Ban, D. Work, J. Herrera, A.M. Bayen, **M. Annavaram**, and Q. Jacobson. Virtual trip lines for distributed privacy-preserving traffic monitoring. In *Proceedings of the 6th international Conference on Mobile Systems, Applications, and Services*, pages 15-28, June, 2008. (Acceptance Rate: 22/123, 18%)
79. **[HPCA]** **M. Annavaram**, E. Grochowski, and P. Reed. Implications of Device Timing Variability on Full Chip Timing. In *Proceedings of the 13th International High-Performance Computer Architecture*, pages 37-45, Feb 2007. (Acceptance rate 32/184, 17%)
80. **[MICRO]** B. Black, **M. Annavaram**, N. Brekelbaum, J. DeVale, L. Jiang, G.H. Loh, D. McCauley, P. Morrow, D.W. Nelson, D. Pantuso, P. Reed, J. Rupley, S. Shankar, J. Shen, and C. Webb. Die Stacking (3D) Microarchitecture. In *Proceedings of the 39th International Symposium on Microarchitecture*, pages 469-479, Dec 2006 (Acceptance Rate :42/174, 24%). (B.Black and M. Annavaram are senior authors with equal contributions)
81. **[ISCA]** **M. Annavaram**, E. Grochowski, J. Shen. Mitigating Amdahl's Law Through EPI Throttling. In *Proceedings of the 32nd International Symposium on Computer Architecture*, pages 298-309, June 2005 (Acceptance Rate 45/194, 23%).
82. **[MICRO]** **M. Annavaram**, R. Rakvic, M. Polito, J. Bouguet, R. Hankins, B. Davies. The Fuzzy Correlation between Code and Performance Predictability. In *Proceedings of the 37th International Symposium on Microarchitecture*, pages 93-104, Dec 2004 (Acceptance Rate: 29/158, 18%).
83. **[MICRO]** R. Hankins*, T. Diep, **M. Annavaram**, B. Hirano, H. Eri, H. Nueckel, and J. Shen. Scaling and Characterizing Database Workloads: Bridging the Gap between Research and Practice. In *Proceedings of the 36th International Symposium on Microarchitecture*, pages 151-162, Dec 2003 (Acceptance Rate: 35/134, 26%)

84. [ICCD] **M. Annavaram**, T. Diep and J. Shen. Branch Behavior of a Commercial OLTP Workload on Intel IA32 Processors. In *Proceedings of the International Conference on Computer Design*, pages 242-248, Sept 2002.
85. [ISCA] **M. Annavaram**, J. M. Patel and E. S. Davidson. Data Prefetching by Dependence Graph Precomputation. In *Proceeding of the 28th International Symposium on Computer Architecture*, pages 52-61, July 2001 (Acceptance Rate 24/163, 15%). (E.S. Davidson is Annavaram's PhD Thesis advisor)
86. [HPCA] **M. Annavaram**, J. M. Patel and E. S. Davidson. Call Graph Prefetching for Database Applications. In *Proceedings of the Seventh International High-Performance Computer Architecture*, pages 281-290, Jan 2001 (Acceptance Rate: 26/110, 24%)
87. [ISPASS] **M. Annavaram**, G. S. Tyson and E. S. Davidson. Instruction Overhead and Data Locality Effects in Superscalar Processors. In *Proceedings of the International Symposium on Performance Analysis of Systems and Software*, pages 95-100, April 2000.
88. [IPDPS] **M. Annavaram**, W. Najjar. Comparison of Two Storage Models in Data-Driven Multithreaded Architectures. In *proceedings of the Symposium on Parallel Distributed Computing*, pages 100-110, October 1996. (W. Najjar is MS Thesis advisor)

Journal Papers

1. [IEEE TAS] C. Fourie et. Al., 2023. Results from the ColdFlux Superconductor Integrated Circuit Design Tool Project. *IEEE Transactions on Applied Superconductivity*.
2. [Optics] A. Minoofar, N. Karapetyan, A. Almainan, H. Zhou, H. Song, K. Zou, W. Ko, M. Ramakrishnan, M. Annavaram, J.L. Habif, M. Tur, and A.E. Willner, Demonstration of optically-assisted reconfigurable average of two 20-Gbaud 4-phase-encoded data channels using nonlinear wave mixing. *Opt. Lett.* 48, 4617-4620 (2023).
3. [IEEE TAS] M. Pedram, M. Annavaram et al., ColdFlux Superconducting EDA and TCAD Tools Project: Overview and Progress. *IEEE Transactions on Applied Superconductivity*, Dec 2018.
4. [IEEE VLSI] K. Matam*, M. Abdel-Majeed*, M. Annavaram. Efficient automatic parallelization of a single GPU program for a multiple GPU system. *IEEE Journal of VLS-Integration*. Dec 2018.
5. [IEEE TOC] Y. Oh, K. Kim, M. Yoon, J. Park, Y. Park, M. Annavaram, W. Ro. Adaptive Cooperation of Prefetching and Warp Scheduling on GPUs. *IEEE Transactions on Computer*, Oct 2018
6. [IEEE TOC] S. Lee, K. Kim, G. Koo*, H. Jeon*, **M. Annavaram** and W. Ro. Improving Energy Efficiency of GPUs through Data Compression and Compressed Execution. *IEEE Transactions on Computers*, Vol. 66, no. 5, Oct 2016.

7. [IEEE TOC] M. Manocheri*, **M. Annavaram** M. Dubois. CPCC: Correctable Parity Protected Cache. *IEEE Transactions on Computers*, Vol.63, no. 10, Oct 2014.
8. [MICRO TOPPICKS] D. Wong* and **M. Annavaram**. Scaling the energy proportionality wall through server-level heterogeneity. In *IEEE MICRO TOPPICKS*, Vol 33, no.3, June 2013.
9. [IEEE TOC] K. Patel*, **M. Annavaram** and **M. Pedram**. NFRA: Generalized Network Flow Based Resource Allocation for Hosting Centers. In *IEEE Transactions on Computers*, Vol 62, no.9, 2013.
10. [COMM] U. Mitra, A. Emken, S. Lee, M. Li, V. Rozgic, G. Thatte, H. Vathsangam, D. Zois, M. Levorato, **M. Annavaram**, S. Narayanan and D. Spruijt-Metz, G Sukhatme. KNOWME: A Case Study in Wireless Body Area Sensor Network Design. In *IEEE Communication Magazine*, Vol 50, no. 5, March 2012.
11. [JPAH] A. Emken, M. Li, G. Thatte, S. Lee*, **M. Annavaram**, U. Mitra, S. Narayanan and D. Spruijt-Metz. Recognition of Physical Activities in Overweight Hispanic Youth Using KNOWME Networks. In *Journal of Physical Activity and Health*, Vol.9, no.3, March 2012.
12. [ToMC] B. Hoh, M. Gruteser, R. Herring, J. Ban, D. Work, J. Herrera, A.M. Bayen, **M. Annavaram**, and Q. Jacobson. Enhancing Privacy and Accuracy in Probe Vehicle-Based Traffic Monitoring via Virtual Trip Lines. In *IEEE Transactions on Mobile Computing*, Vol. 11, no.5, May 2012.
13. [TECS] G. Thatte, M. Li, S. Lee*, A. Emken, S. Narayanan, U. Mitra, D. Spruijt-Metz and **M. Annavaram**. KNOWME: An Energy-Efficient and Multimodal Body Area Sensing System for Physical Activity Monitoring. *ACM Transactions in Embedded Computing Systems (TECS)*, Special Issue on Wireless Health Systems. Vol. 11, no. S2, August 2012.
14. [TSP] G. Thatte, M. Li, S. Lee*, A. Emken, **M. Annavaram**, S. Narayanan, D. Spruijt-Metz and U. Mitra. Optimal Time-Resource Allocation for Energy-Efficient Physical Activity Detection. *IEEE Transactions on Signal Processing*, Vol. 59, no.4, April 2011.
15. [LNCS] N. Madan*, A. Buyuktosunoglu, P. Bose, and, **M. Annavaram**. Guarded Power Gating in a Multi-Core Setting. *Lecture Notes in Computer Science*, Springer Publication, Vol 6162, Feb 2011.
16. [LNCS] S. Ghosh*, M. Redekopp and **M. Annavaram**. KnightShift: Shifting the I/O Burden in Datacenters to Management Processor for Energy Efficiency. *Lecture Notes in Computer Science*, Springer Publication, Vol 6161, 2010.
17. [ToNSRE] M. Li, V. Rozgic, G. Thatte, S. Lee*, A. Emken, **M. Annavaram**, U. Mitra, D. Spruijt-Metz and S. Narayanan. Multimodal Physical Activity Recognition by Fusing Temporal and Cepstral Information. *Transactions on Neural Systems & Rehabilitation Engineering*, Vol. 18, no.4, 2010.
18. [CAN] J. Chen, **M. Annavaram**, and M. Dubois. SlackSim: a platform for parallel simulations of CMPs on CMPs. *SIGARCH Computer Architecture News*, Vol. 37, no.2, July 2009.

19. [TOCS] **M. Annavaram**, J. M. Patel and E. S. Davidson. Call Graph Prefetching for Database Applications. *ACM Transactions on Computer Systems*, Vol. 21, no.4, Nov 2003.

Workshop and Invited Papers

1. [ReML] L. Gao*, Y. Niu, T. Tang*, S. Avestimehr, **M. Annavaram**. Ethos: Rectifying Language Models in Orthogonal Parameter Space. *ReML workshop@AAAI 2024*, Feb 2024.
2. [DPML] H. Hashemi*, Y. Wang*, **M. Annavaram**. Privacy and Integrity Preserving Training Using Trusted Hardware, *DPML workshop@ICLR 2021*, May 2021.
3. [SSML] H. Hashemi*, Y. Wang*, C. Guo, **M. Annavaram**. Byzantine-Robust and Privacy-Preserving Framework for FedML, *SSML workshop@ICLR 2021*, May 2021.
4. [DPML] C. He, K. Balasubramanian, E. Ceyani, Y. Ron, P. Zhao, J. Huang, **M. Annavaram, and S. Avestimehr**. FedGraphNN: A Federated Learning System and Benchmark for Graph Neural Networks, *DPML workshop@ICLR 2021*, May 2021.
5. [GNNSys] C. He, K. Balasubramanian, E. Ceyani, Y. Ron, P. Zhao, J. Huang, **M. Annavaram, and S. Avestimehr**. FedGraphNN: A Federated Learning System and Benchmark for Graph Neural Networks, *GNNSys workshop@MLSys 2021*, March 2021.
6. [SPICYFL] C. He, S. Li, ..., **M. Annavaram, and S. Avestimehr**. FedML: A Research Library and Benchmark for Federated Machine Learning, *SpicyFL workshop*, Dec 2020.
7. [CWWMCA] H. Hashemi, and **M. Annavaram**. DarKnight: Privacy protected machine learning, *6th Career Workshop for Women and Minorities in Computer Architecture*, Oct 2020.
8. [CWWMCA] T. Tan, and **M. Annavaram**. Scalability of coded machine learning, *6th Career Workshop for Women and Minorities in Computer Architecture*, Oct 2020.
9. [NAS] C. He, **M. Annavaram, and S. Avestimehr**. Towards non-IID and invisible data with FedNAS: Federated Deep Learning via Neural Architecture Search, *Workshop on Neural Architecture Search and Beyond for Representation Learning*, Aug 2020.
10. [EUCAS] N. Katam, H. Zha, M. Pedram and **M. Annavaram**. Steroid: A High Performance SFQ Architecture Using Dual-bit Boolean Logic, In *Proceedings of the 14th European Conference on Applied Superconductivity (EUCAS)*, Sept 2019.
11. [EUCAS] N. Katam, H. Zha, M. Pedram and **M. Annavaram**. High Capacity Fluxon Storage for Microprocessor Design. In *Proceedings of the 14th European Conference on Applied Superconductivity (EUCAS)*, Sept 2019.
12. [ISEC] **M. Annavaram**, P. Beerel, S. Gupta, S. Nazarem, and M. Pedram. Progress Towards an Open-Source Front End CAD Flow for RSFQ. In

- proceedings of 17th International Superconductive Electronics Conference*, July 2019.
13. [CodML] K. Narra, Z. Lin, G. Ananthanarayana, S. Avestimeher, and M. Annavaram. Collage-CNN: Using Coded Redundancy for Low Variance Distributed Image Classification, In *Proceedings of the Coded Machine Learning Workshop (CodML)*, July 2019.
 14. [GPUDEPEND] W. Dweik*, M. Abdel-Majeed* and M. Annavaram. Tolerating Hard Faults in GPGPUs. In the *International Workshop on Dependable GPU Computing*, March 2014.
 15. [WEED] D. Wong* and M. Annavaram. Evaluating A Prototype KnightShift-enabled Server. In the *Workshop on Energy Efficient Design*, June 2012.
 16. [MOBS] J.Chen, K. Dabbiru*, M. Annavaram and M. Dubois. Adaptive and Speculative Slack Simulations of CMPs on CMPs. In the Sixth Annual Workshop on Modeling, Benchmarking and Simulation (*MOBS*), June 2010.
 17. [MOBS] J. Suh*, M. Annavaram and M. Dubois. Soft Error Benchmarking for L2 Cache with PARMA. In the Sixth Annual Workshop on Modeling, Benchmarking and Simulation (*MOBS*), June 2010.
 18. [WEED] S. Ghosh*, M. Redekopp and M. Annavaram. KnightShift: Shifting the I/O Burden in Datacenters to Management Processor for Energy Efficiency. In the *Workshop on Energy Efficient Design*, June 2010.
 19. [WEED] N. Madan*, A. Buyuktosunoglu, P. Bose, and, M. Annavaram. Guarded Power Gating in a Multi-Core Setting. In the *Workshop on Energy Efficient Design*, June 2010.
 20. [CENS] G. Thatte, M. Li, A. Emken, U. Mitra, S. Narayanan, M. Annavaram and D. Spruijt-Metz. Energy-Efficient Activity-Detection via Multihypothesis Testing for Pediatric Obesity. In the *7th Annual CENS Research Review*, October 2009.
 21. [mHealth] D. Spruijt-Metz, S. Narayanan, U. Mitra, G. Sukhatme, M. Li, G. Thatte, A. Emken, S. Lee, H. Vathsangam and M. Annavaram. KNOWME Networks: Mobile Device Biomonitoring to Prevent and Treat Obesity in Underserved Minority Youth. In the *mHealth Summit*, October 2009.
 22. [BIC] S. Lee*, M. Annavaram, G. Thatte, V. Rozgic, M.Li, U. Mitra, S. Narayanan, A. Emken, and D. Spruijt-Metz. Sensing for Obesity: KNOWME Implementation and Lessons for an Architect. In *Workshop on Biomedicine in Computing: Systems, Architectures, and Circuits (BiC)*, June 2009.
 23. [ICDAM] D. Spruijt-Metz, M. Li, G. Thatte, G. Sukhatme, M. Annavaram, S. Ghosh, V. Rozgic, U. Mitra, N. Medvidovic, B. Belcher, and S. Narayanan. Differentiating physical activity modalities in youth using heartbeat waveform shape and differences between adjacent waveforms. In *proceedings of the 7th International Conference on Diet and Activity Methods (ICDAM 7)*, June 2009.
 24. [BodyNets] G. Thatte, V. Rozgic, M. Li, S. Ghosh*, U. Mitra, S. Narayanan, M. Annavaram, and D. Spruijt-Metz. Optimal time-resource allocation for activity-detection via multimodal sensing. In *proceedings of the Fourth International Conference on Body Area Networks*, April 2009.
 25. [URBANSENSE] M. Annavaram, N. Medvidovic, U. Mitra, S. Narayanan, G. Sukhatme, Z. Meng, S. Qiu, R. Kumar, G. Thatte, and D. Spruijt-Metz.

- Multimodal sensing for pediatric obesity applications. In *Proceedings of UrbanSense08*, November 2008.
26. [MODUS] **M. Annavaram**, Q. Jacobson. HangOut: A Privacy Preserving Location Based Social Networking Service. In *Workshop on Mobile Devices and Urban Sensing, (Invited Paper)* April 2008.
 27. [ISPD] **M. Annavaram**, E. Grochowski, and P. Reed. Implications of Device Timing Variability on Full Chip Timing. In *Proceedings of the International Symposium on Physical Design, (Invited Paper)* April 2008.
 28. [CAECW] R. Hankins*, **M. Annavaram**, B. Hirano, J. Patel and J. Shen. Comparing OLTP Scaling Behavior on Intel® Xeon™ and Itanium® 2 Processors. In *the seventh Workshop on Computer Architecture Evaluation using Commercial Workloads*, Feb 2004.
 29. [MEM] P. Kundu, **M. Annavaram**, T. Diep and J. Shen. A Case for Shared Instruction Cache on Chip Multiprocessors Running OLTP. In *the Memory Performance: Dealing With Applications, Systems and Architectures Workshop*, Sep 2003.
 30. [CAECW] **T. Diep, M. Annavaram**, B. Hirano and J. Shen. Analyzing Performance Characteristics of OLTP Cached Workloads by Linear Interpolation. In *the sixth Workshop on Computer Architecture Evaluation using Commercial Workloads*, Feb 2003.
 31. [WWC] J. Rupley II, **M. Annavaram**, J. DeVale, T. Diep and B. Black. Comparing and Contrasting a Commercial OLTP Workload with CPU2000 on IPF. In *the fifth Workshop on Workload Characterization*, Nov 2002.
 32. [EPIC] R. Rakvic, E. Grochowski, B. Black, **M. Annavaram**, T. Diep and J. Shen. Performance Advantage of the Register Stack in Intel Itanium Processors. In *the second Workshop on Explicitly Parallel Instruction Computing Architecture and Compilers*, Nov 2002.

Patents

1. US-2020310690-A1, *Dynamic near-data processing control mechanism based on computer resource availability on solid-state disk platforms*, Murali Annavaram, Gunjae Koo, Kiran Kumar Matam, Hung-Wei Tseng. Filed 10/2017.
2. US-9008735-B2, *Runtime selection of most energy-efficient approach for services requested by mobile applications*, Murali Annavaram, Sangwon Lee. Issued 4/2015.
3. US-7437581-B2, *Method and apparatus for varying energy per instruction according to the amount of available parallelism*, Edward Grochowski, John Shen, Hong Wang, Doron Orenstein, Gad S Sheaffer, Ronny Ronen, Murali Annavaram. Issued 10/2008.
4. US-2012131366-A1, *Load balancing for multi-threaded applications via asymmetric power throttling*, Ryan Rakvic, Hankins Richard A, Ed Grochowski, Hong Wang, Murali Annavaram, Poulsen David K, Sanjiv Shah, John Shen, Gautham China. Issued 9/2014.

5. US-8631290-B2, *Automated detection of and compensation for guardband degradation during operation of clocked data processing circuit*, Bardia Zandian, Murali Annavaram. Issued 1/2014.
6. US-2010052730-A1, *Method and apparatus for late timing transition detection*, Edward Grochowski, Chris Wilkerson, Lu Shih-Lien L, Murali Annavaram. Issued 2/2012.
7. HK-1121257-A1, *Prefetching from a dynamic random access memory to a static random access memory*, Bryan P Black, Murali Annavaram, Donald W Mccauley, John P Devale. Issued 10/2011.
8. WO-2009068970-A1, *Methods, apparatuses, and computer program products for traffic data aggregation using virtual trip lines and a combination of location and time based measurement triggers in gps-enabled mobile handsets*, Jacobson Quinn, Hoh Baik, Murali Annavaram, David Sutter. Issued 3/2011.
9. US-2009143966-A1, *Methods, apparatuses, and computer program product for traffic data aggregation using virtual trip lines and gps-enabled mobile handsets*, Quinn Jacobson, Baik Hoh, Murali Annavaram. Issued 10/2010.
10. US-7480838-B1, *Method, system and apparatus for detecting and recovering from timing errors*, Chris Wilkerson, Shih-Lien L. Lu, Edward Grochowski, Murali Annavaram. Issued 1/2009.
11. US-2007220207-A1, *Transferring data from stacked memory*, Bryan Black, Murali Annavaram, Paul Reed. Issued 9/2007.
12. US-2004049666-A1 *Method and apparatus for variable pop hardware return address stack*, Murali Annavaram, Trung Diep, John Shen. Issued 3/2004.

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

ARM LTD.,

Plaintiff,

v.

QUALCOMM INC., QUALCOMM
TECHNOLOGIES, INC. and NUVIA, INC.,

Defendants

C.A. No. 22-1146 (MN)

OPENING EXPERT REPORT OF DR. MURALI ANNAVARAM

APPENDIX

B

EXPERT REPORT OF MURALI ANNAVARAM
HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

APPENDIX B

MATERIALS CONSIDERED

1. Documents referenced in my Opening Expert Report.
2. Discovery Responses, witness deposition transcripts, and exhibits thereto, including at least:

2023.09.22 - Singh, Rohit Deposition Transcript
2023.10.12 - Gulati, Manu Deposition Transcript
2023.10.27.23 - Sharma, Nitin Rough Deposition Transcript
2023.11.03 - Williams III, Gerard Rough Deposition Transcript
2023.11.08 - Ashgar, Ziad Rough Deposition Transcript
2023.10.26 – Defendants’ Supplemental and Amended Response and Objection to Plaintiff’s First Set of Interrogatories (No.5)
2023.11.17 – Defendants’ Supplemental Responses and Objections to Plaintiff’s First Set of Interrogatories (Nos. 7-12)
2023.11.29 - Bos, Lynn Rough Deposition Transcript
2023.12.08 - Balakrishnan, Geeta Rough Deposition Transcript
Bos Dep. Exhibit 16
Williams Dep. Exhibit 24

3. Scripts, output files, and tracking inventories relating to ARM RTL, including at least:

QCARM_3972000
QCARM_7482165
QCARM_7482214
QCARM_7482219
QCARM_7530238
QCARM_7531465
QCARM_7530244
QCARM_7530242
QCARM_7530240
QCARM_7530239

EXPERT REPORT OF MURALI ANNAVARAM
HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

QCARM_7530246

4. Produced documents and websites related to the formation and history of Nuvia, including at least:

QCARM_0002749

QCARM_2414840

QCARM_3427968

https://www.qualcomm.com/news/releases/2021/03/qualcomm-completes-acquisition-nuvia

5. Produced documents and websites related to the ARM Connect platform, including at least:

QCARM_3616548 (ARM Connect User Guide)
--

https://developer.arm.com/documentation/107571/latest/ (ARM Connect and Product Download Hub Key Differences v2.0)
--

https://developer.arm.com/documentation/107572/latest/ (ARM Product Download Hub Getting Started Guide v4.0)
--

6. Websites related to Git and Gerrit tools, including at least:

https://www.gerritcodereview.com/about.html

https://www.gerritcodereview.com/#serve-git

https://git-scm.com/book/en/v2/Getting-Started-What-is-Git%3F

https://en.wikipedia.org/wiki/SHA-1

7. Source code including source code comparisons, including at least:

[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]

EXPERT REPORT OF MURALI ANNAVARAM
HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]

8. Interviews with Qualcomm engineers from the development teams, including at least:

December 11, 2023 Bob Pflederer Interview
December 12, 2023 Bob Plederer and Nick Jones Interview
December 13, 2023 Nick Jones Interview
December 20, 2023 Bob Plederer and Nick Jones Interview

9. Git logs, including at least:

[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]

EXPERT REPORT OF MURALI ANNAVARAM
HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[illegible]

EXPERT REPORT OF MURALI ANNAVARAM
HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[illegible]

EXPERT REPORT OF MURALI ANNAVARAM
HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[illegible]

EXPERT REPORT OF MURALI ANNAVARAM
HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

EXPERT REPORT OF MURALI ANNAVARAM
HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[illegible]

EXPERT REPORT OF MURALI ANNAVARAM
HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[illegible]

EXPERT REPORT OF MURALI ANNAVARAM
HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[illegible]

EXPERT REPORT OF MURALI ANNAVARAM
HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[illegible]

EXPERT REPORT OF MURALI ANNAVARAM
HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[illegible]

EXPERT REPORT OF MURALI ANNAVARAM
HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY

[illegible]

**EXPERT REPORT OF MURALI ANNAVARAM
HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEY'S EYES ONLY**

[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]

10. Scripts for operating on code repositories, including at least:

[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]

11. Qualcomm source code snapshots produced in this case available on a source code review machine at the ProSearch source code facility in Los Angeles, California.

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

ARM LTD.,

Plaintiff,

v.

QUALCOMM INC., QUALCOMM
TECHNOLOGIES, INC. and NUVIA, INC.,

Defendants


C.A. No. 22-1146 (MN)

OPENING EXPERT REPORT OF DR. MURALI ANNAVARAM

APPENDIX

C

Appendix C contains the following documents, which have been designated as “HIGHLY CONFIDENTIAL - SOURCE CODE - ATTORNEYS’ EYES ONLY” pursuant to the Protective Order in this action.



Pursuant to the Protective Order, a printed copy of the above-listed documents is being sent via FedEx First Overnight Delivery, Tracking Number 7745 6587 5859 to:

Fahd Hussein Patel, Esquire
Morrison & Foerster LLP
2100 L Street, NW
Suite 900
Washington, D.C., 20037
Phone: (202) 887-1500